JUNIPER NETWORKS

Junos® OS

Ethernet Interfaces Feature Guide for Routing Devices

Modified: 2018-06-25

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

Junos® OS Ethernet Interfaces Feature Guide for Routing Devices

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

Table of Contents

List of Figures

List of Tables

About the Documentation

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 https://www.juniper.net/documentation/.

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 https://www.juniper.net/books.

Supported Platforms

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

- ACX Series
- M Series
- MX Series
- T Series
- PTX Series

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 xlix 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 l defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
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.	

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
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**.
> (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 https://www.juniper.net/documentation/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at https://www.juniper.net/documentation/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 https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: https://www.juniper.net/cm/

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

Opening a Case with JTAC

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

- Use the Case Management tool in the CSC at https://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 https://www.juniper.net/support/requesting-support.html.

PART 1

Ethernet Interfaces

CHAPTER 1

Ethernet Interfaces Overview

Ethernet Interfaces Overview

Ethernet was developed in the early 1970s at the Xerox Palo Alto Research Center (PARC) as a data-link control layer protocol for interconnecting computers. It was first widely used at 10 megabits per second (Mbps) over coaxial cables and later over unshielded twisted pairs using 10Base-T. More recently, 100Base-TX (Fast Ethernet, 100 Mbps), Gigabit Ethernet (1 gigabit per second [Gbps]), 10-Gigabit Ethernet (10 Gbps), and 100-Gigabit Ethernet (100 Gbps) have become available.

Juniper Networks routers support the following types of Ethernet interfaces:

- Fast Ethernet
- Tri-Rate Ethernet copper
- Gigabit Ethernet
- Gigabit Ethernet intelligent queuing (IQ)
- Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet
- 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM)
- 100-Gigabit Ethernet
- Management Ethernet interface, which is an out-of-band management interface within the router
- Internal Ethernet interface, which connects the Routing Engine to the packet forwarding components
- Aggregated Ethernet interface, a logical linkage of Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet physical connections

Related Documentation

- Configuring Ethernet Physical Interface Properties on page 6
- MX Series Router Interface Identifiers on page 4
- Configuring MAC Address Filtering for Ethernet Interfaces on page 14
- Configuring Ethernet Loopback Capability on page 18
- Configuring Flow Control on page 12
- Ignoring Layer 3 Incomplete Errors on page 19
- Configuring the Link Characteristics on Ethernet Interfaces on page 13
- Configuring Gratuitous ARP on page 19
- Adjusting the ARP Aging Timer on page 21
- Configuring the Interface Speed on Ethernet Interfaces on page 8
- Configuring the Ingress Rate Limit on page 10
- Configuring Multicast Statistics Collection on Ethernet Interfaces on page 22
- Configuring Weighted Random Early Detection on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

MX Series Router Interface Identifiers

Juniper Networks MX Series 3D Universal Edge Routers support several types of line cards, including Dense Port Concentrators (DPCs), Flexible Port Concentrators (FPCs) with associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) with associated Modular Interface Cards (MICs), or MICs. FPCs are populated with PICs for various interface types. DPCs and MPCs with associated MICs, and MICs support a variety of port configurations and combine the functions of FPCs and the PICs. The configuration syntax for each type of line card is the same: *type-fpc/pic/port*.

Ports are numbered from 0 through 9 for Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. Port numbers are always 0 for 10-Gigabit Ethernet interfaces.

NOTE: In certain displays, the MX Series routers identify the Packet Forwarding Engine (PFE) rather than the PIC number. PFE 0 corresponds to PIC 0, PFE 1 corresponds to PIC 2, PFE 2 corresponds to PIC 1, and PFE 3 corresponds to PIC 3.

Related Documentation

- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

CHAPTER 2

Performing Initial Configuration for Ethernet Interfaces

Example: Configuring Fast Ethernet Interfaces

The following configuration is sufficient to get a Fast Ethernet interface up and running. By default, IPv4 Fast Ethernet interfaces use Ethernet version 2 encapsulation.

```
[edit]
user@host# set interfaces fe-5/2/1 unit 0 family inet address local-address
user@host# show
interfaces {
```

```
    fe-5/2/1 {
      unit 0 {
        family inet {
          address local-address;
        }
      }
    }
  }
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring Gigabit Ethernet Interfaces

The following configuration is sufficient to get a Gigabit Ethernet, Tri-Rate Ethernet copper, or 10-Gigabit Ethernet interface up and running. By default, IPv4 Gigabit Ethernet interfaces on MX Series, M Series, and T Series routers use 802.3 encapsulation.

```
[edit]
user@host# set interfaces ge-2/0/1 unit 0 family inet address local-address
user@host# show
interfaces {
  ge-2/0/1 {
    unit 0 {
      family inet {
        address local-address;
      }
    }
  }
}
```

The M160, M320, M120, T320, and T640 2-port Gigabit Ethernet PIC supports two independent Gigabit Ethernet links.

Each of the two interfaces on the PIC is named:

ge-*fpc/pic/*[0.1]

Each of these interfaces has functionality identical to the Gigabit Ethernet interface supported on the single-port PIC.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Ethernet Physical Interface Properties

1. To configure Fast Ethernet-specific physical interface properties, include the **fastether-options** statement at the **[edit interfaces fe-*fpc/pic/port*]** hierarchy level:

   ```
   [edit interfaces fe-fpc/pic/port]
   user@host# set fastether-options;
   ```

NOTE: The **speed** statement applies to the management Ethernet interface (**fxp0** or **em0**), the Fast Ethernet 12-port and 48-port Physical Interface Card (PIC) interfaces and the MX Series Tri-Rate Ethernet copper interfaces. The Fast Ethernet, **fxp0**, and **em0** interfaces can be configured for 10 Mbps or 100 Mbps (**10m | 100m**). The MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps (**10m | 100m | 1g**). The 4-port and 8-port Fast Ethernet PICs support a speed of 100 Mbps only.

MX Series routers support Gigabit Ethernet automatic line sensing of MDI (Media Dependent Interface) and MDIX (Media Dependent Interface with Crossover) port connections. MDI is the Ethernet port connection typically used on network interface cards (NIC). MDIX is the standard Ethernet port wiring for hubs and switches. This feature allows MX Series routers to automatically detect MDI and MDIX connections and configure the router port accordingly. You can disable this feature by using the **no-auto-mdix** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level.

NOTE: Junos OS supports Ethernet host addresses with no subnets. This enables you to configure an Ethernet interface as a host address (that is, with a network mask of /32), without requiring a subnet. Such interfaces can serve as OSPF point-to-point interfaces, and MPLS is also supported.

2. To configure physical interface properties specific to Gigabit Ethernet and 10-Gigabit Ethernet, include the **gigether-options** statement at the **[edit interfaces ge-*fpc/pic/port*]** or **[edit interfaces xe-*fpc/pic/port*]** hierarchy level:

   ```
   [edit interfaces ge-fpc/pic/port]
   user@host# set gigether-options ;
   ```

3. For 10-Gigabit Ethernet DWDM-specific physical interface properties, include the **optics-options** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level:

   ```
   [edit interfaces ge-fpc/pic/port]
   user@ host# set optics-options;
   ```

   To configure Gigabit Ethernet IQ-specific physical interface properties, include the **gigether-options** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level. These statements are supported on 10-Gigabit Ethernet IQ2 and IQ2-E PIC. Some of these statements are also supported on Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

   ```
   [edit interfaces ge-fpc/pic/port]
   user@host# set gigether-options {
   ```

4. To configure 10-Gigabit Ethernet physical interface properties, include the **lan-phy** or **wan-phy** statement at the **[edit interfaces xe-*fpc/pic/port* framing]** hierarchy level.

   [edit interfaces *interface-name*]
   user@host# set framing;

5. To configure OAM 802.3ah support for Ethernet interfaces, include the **oam** statement at the **[edit protocols]** hierarchy level.

   [edit protocols]
   user@host# set oam;

6. To configure Gigabit Ethernet IQ-specific logical interface properties, include the **input-vlan-map**, **output-vlan-map**, **layer2-policer**, and **vlan-tags** statements at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**.

   [edit interfaces *interface-name* unit *logical-unit-number*]
   user#host# set input-vlan-map;
   user@host# set output-vlan-map;
   user#host# set layer2-policer{
   user@host# set vlan-tags inner *tpid.vlan-id* outer *tpid.vlan-id*;

7. To configure aggregated Ethernet-specific physical interface properties, include the **aggregated-ether-options** statement at the **[edit interfaces ae*x*]** hierarchy level:

   [edit interfaces ae*x*]
   user@host# set aggregated-ether-options;

Related Documentation

- 10-Gigabit Ethernet Framing Overview on page 427
- Example: Configuring Gigabit Ethernet Interfaces on page 6
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the Interface Speed on Ethernet Interfaces

For M Series and T Series Fast Ethernet 12-port and 48-port PIC interfaces, the management Ethernet interface (**fxp0** or **em0**), and the MX Series Tri-Rate Ethernet copper interfaces, you can explicitly set the interface speed. The Fast Ethernet, **fxp0**, and **em0** interfaces can be configured for 10 Mbps or 100 Mbps **(10m | 100m)**. The MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps **(10m | 100m | 1g)**. For information about management Ethernet interfaces and to determine the management Ethernet interface type for your router, see *Understanding Management Ethernet Interfaces* and *Supported Routing Engines by Router*MX Series routers, with MX-DPC and Tri-Rate Copper SFPs, support 20x1 Copper to provide backwards compatibility with 100/10BASE-T and 1000BASE-T operation through an Serial Gigabit Media Independent Interface (SGMII) interface.

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

   ```
   [edit ]
   user@host# edit interfaces interface-name
   ```

2. To configure the speed, include the **speed** statement at the **[edit interfaces *interface-name*]** hierarchy level.

   ```
   [edit interfaces interface-name]
   user@host# set speed (10m | 100m | 1g | auto | auto-10m-100m);
   ```

NOTE:

- By default, the M Series and T Series routers management Ethernet interface autonegotiates whether to operate at 10 megabits per second (Mbps) or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the **no-concatenate** statement in the **[edit chassis]** configuration hierarchy.
- Starting with Junos OS Release 14.2 the **auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m**maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.,
- When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.
- If the link partner does not support autonegotiation, configure either Fast Ethernet port manually to match its link partner's speed and link mode. When the link mode is configured, autonegotiation is disabled.
- On MX Series routers with tri-rate copper SFP interfaces, if the port speed is negotiated to the configured value and the negotiated speed and interface speed do not match, the link will not be brought up.
- When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
- Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the **speed** statement, you must include the **link-mode full-duplex** statement at the same hierarchy level.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2 the **auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m** maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.
11.4	Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the **speed** statement, you must include the **link-mode full-duplex** statement at the same hierarchy level.

Related Documentation

- speed on page 1367
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the Ingress Rate Limit

On Fast Ethernet 8-port, 12-port, and 48-port PIC interfaces only, you can apply port-based rate limiting to the ingress traffic that arrives at the PIC.

To configure an ingress rate limit on a Fast Ethernet 8-port, 12-port, or 48-port PIC interface, include the **ingress-rate-limit** statement at the **[edit interfaces *interface-name* fastether-options]** hierarchy level:

[edit interfaces *interface-name* fastether-options]
ingress-rate-limit *rate*;

rate can range in value from 1 through 100 Mbps.

Related Documentation

- ingress-rate-limit on page 1196
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding Flow Control

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

The MX, T, and PTX Series routers support IEEE 802.3X Ethernet PAUSE method of flow control.

• IEEE 802.3X Ethernet PAUSE on page 11

IEEE 802.3X Ethernet PAUSE

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

  In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level.

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

Related Documentation

- **flow-control on page 1170**
- Configuring Flow Control on page 12

Configuring Flow Control

By default, the router or switch imposes flow control to regulate the amount of traffic sent out on a Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interface. Flow control is not supported on the 4-port Fast Ethernet PIC. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the router or switch to permit unrestricted traffic. To disable flow control, include the **no-flow-control** statement:

```
no-flow-control;
```

To explicitly reinstate flow control, include the **flow-control** statement:

```
flow-control;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Related Documentation

- **flow-control on page 1170**
- Ethernet Interfaces Overview on page 3
- *Interfaces Overview for Switches*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the Link Characteristics on Ethernet Interfaces

Full-duplex communication means that both ends of the communication can send and receive signals at the same time. *Half-duplex* is also bidirectional communication, but signals can flow in only one direction at a time.

By default, the router's management Ethernet interface, **fxp0** or **em0**, autonegotiates whether to operate in full-duplex or half-duplex mode. Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet and 10-Gigabit Ethernet, the link partner must also be set to full duplex.

NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is **fxp0**. For T1600 and T4000 routers configured in a routing matrix, and TX Matrix Plus routers, the management Ethernet interface is **em0**.

NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the **fxp0** management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the **fxp0** management Ethernet interface so that the commands reference the **em0** management Ethernet interface instead.

NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

NOTE: Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.

To explicitly configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the **link-mode** statement at the **[edit interfaces *interface-name*]** hierarchy level:

[edit interfaces *interface-name*]

link-mode (full-duplex | half-duplex);

NOTE: Starting with Junos OS release 17.4R1 and later, the **link-mode** configuration is not supported for 10-Gigabit Ethernet interfaces.

Related Documentation

- link-mode on page 1225
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring MAC Address Filtering for Ethernet Interfaces

Enabling Source Address Filtering

On aggregated Ethernet interfaces, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can enable source address filtering to block all incoming packets from a specific MAC address.

To enable the filtering, include the **source-filtering** statement at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

NOTE: When you integrate a standalone T640 router into a routing matrix, the PIC media access control (MAC) addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, when you integrate a T1600 or T4000 router into a routing matrix, the PIC MAC addresses for the integrated T1600 or T4000 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the configuration of a formerly standalone T1600 or T4000 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

When source address filtering is enabled, you can configure the interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the **source-address-filter *mac-address*** statement at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

You can specify the MAC address as ***nn*:*nn*:*nn*:*nn*:*nn*:*nn*** or ***nnnn* .*nnnn*.*nnnn***, where **n** is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include the **source-address-filter** statement multiple times.

NOTE: The **source-address-filter** statement is not supported on Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router); instead, include the **accept-source-mac** statement. For more information, see "Configuring MAC Address Filtering" on page 544.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

Source address filtering does not work when Link Aggregation Control Protocol (LACP) is enabled. This behavior is not applicable to T series routers and PTX Series Packet Transport Routers. For more information about LACP, see "Configuring LACP for Aggregated Ethernet Interfaces" on page 140.

NOTE: On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement at the [edit interfaces **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

NOTE: The **source-address-filter** statement is not supported on MX Series routers with MPC4E (model numbers: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE); instead, include the **accept-source-mac** statement. For more information, see "Configuring MAC Address Filtering" on page 544.

Related Documentation

- source-address-filter on page 1365

- Configuring MAC Address Filtering on page 544
- Configuring LACP for Aggregated Ethernet Interfaces on page 140
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring MAC Address Filtering on PTX Series Packet Transport Routers

This topic describes how to configure MAC filtering on PTX Series Packet Transport Routers. MAC filtering enables you to specify the MAC addresses from which the Ethernet interface can receive packets.

MAC filtering support on PTX Series Packet Transport Routers includes:

- MAC source and destination address filtering for each port.
- MAC source address filtering for each physical interface.
- MAC source address filtering for each logical interface.

When you filter logical and physical interfaces, you can specify up to 1000 MAC source addresses per port.

To configure MAC source address filtering for a physical interface, include the **source-filtering** and **source-address-filter** statements at the **[edit interfaces et-*fpc/pic/port* gigether-options]** hierarchy level:

```
[edit interfaces]
et-x/y/z {
  gigether-options {
    source-filtering;
    source-address-filter {
      mac-address;
    }
  }
}
```

The **source-address-filter** statement configures which MAC source addresses are filtered. The specified physical interface drops all packets from the MAC source addresses you specify. You can specify the MAC address as ***nn:nn:nn:nn:nn:nn*** where ***n*** is a decimal digit. To specify more than one address, include multiple ***mac-address*** options in the **source-address-filter** statement.

To configure MAC source address filtering for a logical interface, include the **accept-source-mac** statement at the **[edit interfaces et-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces]
et-x/y/z {
  gigether-options {
    source-filtering;
  }
  unit logical-unit-number {
```



```
      accept-source-mac {
        mac-address mac-address;
      }
    }
  }
```

The **accept-source-mac** statement configures which MAC source addresses are accepted on the logical interface. You can specify the MAC address as ***nn:nn:nn:nn:nn:nn*** where ***n*** is a decimal digit. To specify more than one address, include multiple **mac-address *mac-address*** options in the **accept-source-mac** statement.

After an interface filter is configured, there is an accounting entry that is associated with the MAC address filter. Counters accumulate if there are packets with matching MAC source addresses. You can use the **show interfaces mac-database** Junos OS CLI command to view the address count.

Related Documentation

- show interfaces mac-database on page 1985

MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces Overview

Junos OS supports the capability to compute MAC address statistics for dynamically learned static and destination MAC addresses on physical interfaces. Starting in Junos OS Release 15.1, Junos OS enables you to configure source MAC (SMAC) address and destination MAC (DMAC) address-based accounting for MAC addresses that are dynamically learned on aggregated Ethernet (ae-) interfaces in routed mode. When you include the **mac-learn-enable** statement at the **[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]** hierarchy level, dynamic learning of source and destination MAC addresses is enabled. By default, this capability is disabled. When dynamic learning of MAC addresses is enabled for AE interfaces in routed mode, the MAC-filter settings are updated for each of the child links of the AE bundle interface. This feature provides for both the configuration of the mac-learn-enable filter and the display of SMAC and DMAC based accounting information on the aggregated interface in the output of the **show interfaces mac-database *interface-name* mac-address *mac-address*** command.

When this functionality is enabled, source and destination MAC addresses-based accounting is supported on the routed interfaces on MX Series routers with DPCs and MPCs. Support for mixed mode LAG interfaces is also available. This feature supports MAC address accounting for AE interfaces in routed mode (for inet family). Destination MAC-based accounting is supported only for MAC addresses dynamically learned at the ingress interface, including each individual child or member link of the AE bundle. This behavior occurs because MPCs do not support destination MAC address learning. As a result, if a packet exits a child link without passing in the ingress direction through that link, destination MAC (DMAC) accounting for this packet occurs at the child link level and this data is not available at the aggregate level. Dynamic learning of MAC addresses can be supported on only the AE interface or on selective individual member links. MAC learning support on the bundle depends on the capability of individual member links. If a link in the bundle does not contain the capability to support MAC learning or accounting, it is disabled on the AE bundle.

The MAC data for the aggregated bundle is displayed by collecting data form individual child links. This data is collected when the command to display the MAC database is triggered from the CLI. This method of data collection implies that based on the number of child links and the size of the MAC database, the time take to display the database differs. This approach to obtain the current snapshot of the MAC database from the currently active child links is used instead of maintaining a database at the Routing Engine because of the dynamic nature of the MAC database and the overhead required to maintain the database information in synchronization with all the child Packet Forwarding Engines. A difference in the DMAC-based accounting for packets generated from the Routing Engine (packets sent in the host path). On DPCs, these packets are accounted in egress direction (Output Packet/Byte count), whereas on MPCs, these packets are not accounted because DMAC learning is not supported. This difference in behavior also occurs between child links on DPCs and MPCs. Because this feature to enable dynamic learning is related to collecting MAC database statistics from child links based on the command issued from the CLI, there is a impact on the time it takes to display the data on the console based on the size of the MAC database and the number of child-links spread across different FPCs. The limit on the maximum number of MAC addresses that can be learned from an interface does not apply to this dynamic learning of MAC addresses functionality.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, Junos OS enables you to configure source MAC (SMAC) address and destination MAC (DMAC) address-based accounting for MAC addresses that are dynamically learned on aggregated Ethernet (ae-) interfaces in routed mode.

Related Documentation

- mac-learn-enable on page 1245

Configuring Ethernet Loopback Capability

By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

NOTE: If you configure a local loopback on a 1-port 10-Gigabit IQ2 and IQ2-E PIC using the **loopback** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level, the transmit-path stops working, causing the remote end to detect a link down.

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

Related Documentation

- loopback on page 1239
- Ethernet Interfaces Overview on page 3
- *Interfaces Overview for Switches*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Ignoring Layer 3 Incomplete Errors

By default, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces count Layer 3 incomplete errors. You can configure the interface to ignore Layer 3 incomplete errors.

To ignore Layer 3 incomplete errors, include the **ignore-l3-incompletes** statement:

```
ignore-l3-incompletes;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

Related Documentation

- ignore-l3-incompletes on page 1192
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests help detect duplicate IP addresses. A gratuitous ARP is a broadcast request for a router's own IP address. If a router or switch sends an ARP request for its own IP address and no ARP replies are received, the router- or switch-assigned IP address is not being used by other nodes.

However, if a router or switch sends an ARP request for its own IP address and an ARP reply is received, the router- or switch-assigned IP address is already being used by another node.

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router or switch receives a gratuitous ARP reply, the router or switch can insert an entry for that reply in the ARP cache. By default, updating the ARP cache on gratuitous ARP replies is disabled on the router or switch.

To enable updating of the ARP cache for gratuitous ARPs:

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Include the **gratuitous-arp-reply** statement.

   ```
   [edit interfaces interface-name]
   user@host# set gratuitous-arp-reply
   ```

To restore the default behavior, that is, to disable updating of the ARP cache for gratuitous ARP, delete the **gratuitous-arp-reply** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete gratuitous-arp-reply;
```

By default, the router or switch responds to gratuitous ARP requests. However, on Ethernet interfaces, you can disable responses to gratuitous ARP requests.

To disable responses to gratuitous ARP requests:

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Include the **no-gratuitous-arp-request** statement.

   ```
   [edit interfaces interface-name]
   user@host# set no-gratuitous-arp-request
   ```

To return to the default—that is, to respond to gratuitous ARP requests—delete the **no-gratuitous-arp-request** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete no-gratuitous-arp-request
```

Related Documentation

- gratuitous-arp-reply on page 1181
- no-gratuitous-arp-request on page 1274
- Ethernet Interfaces Overview on page 3

- *Interfaces Overview for Switches*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Adjusting the ARP Aging Timer

By default, the ARP aging timer is set at 20 minutes. In environments with many directly attached hosts, such as metro Ethernet environments, increasing the amount of time between ARP updates by configuring the ARP aging timer can improve performance in an event where having thousands of clients time out at the same time might impact packet forwarding performance. In environments where there are devices connected with lower ARP aging timers (less than 20 minutes), decreasing the ARP aging timer can improve performance by preventing the flooding of traffic toward next hops with expired ARP entries. In most environments, the default ARP aging timer value does not need to be adjusted.

To configure the system-wide ARP aging timer, include the **aging-timer** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
user@host# aging-timer minutes
```

The aging timer range is from 1 through 240 minutes. The timer value you configure takes effect as ARP entries expire. In other words, each subsequent refreshed ARP entry receives the new timer value. The new timer value does not apply to ARP entries that exist at the time you commit the configuration.

For more information about statements you can configure at the **[edit system]** hierarchy level, see the *Junos OS Administration Library*.

Related Documentation

- arp on page 1082
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Weighted Random Early Detection

On M7i, M10i, M40e, M320, M120, and T Series routers, the Ethernet IQ2 and IQ2-E PIC families extend CoS functionality by supporting network congestion avoidance with weighted random early detection (WRED).

Related Documentation

- For information on configuring WRED, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Multicast Statistics Collection on Ethernet Interfaces

T Series and TX Matrix routers support multicast statistics collection on Ethernet interfaces in both ingress and egress directions. The multicast statistics functionality can be configured on a physical interface thus enabling multicast accounting for all the logical interfaces below the physical interface.

The multicast statistics information is displayed only when the interface is configured with the **multicast-statistics** statement, which is not enabled by default.

Multicast statistics collection requires at least one logical interface is configured with family inet and/or inet6; otherwise, the commit for **multicast-statistics** will fail.

The multicast in/out statistics can be obtained via interfaces statistics query through CLI and via MIB objects through SNMP query.

To configure multicast statistics:

1. Include the **multicast-statistics** statement at the **[edit interfaces interface-name]** hierarchy level.

An example of a multicast statistics configuration for a Ethernet interface follows:

```
[edit interfaces]
  ge-fpc/pic/port {
  multicast-statistics;
}
```

To display multicast statistics, use the **show interfaces** ***interface-name*** **statistics detail** command.

Related Documentation

- *multicast-statistics*
- Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces on page 196
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router

The router internal Ethernet interface connects the Routing Engine with the router's packet forwarding components. The Junos OS automatically configures internal Ethernet interfaces. For TX Matrix Plus routers, the internal Ethernet interfaces are **ixgbe0** and **ixgbe1**. For T1600 routers configured in a routing matrix, the internal Ethernet interfaces are **bcm0** and **em1**. For more information about internal Ethernet interfaces, see *Understanding Internal Ethernet Interfaces*.

NOTE: Do not modify or remove the configuration for the internal Ethernet interface that the Junos OS automatically configures. If you do, the router will stop functioning.

The following example is a sequence of **show interfaces** commands issued in a Junos OS command-line interface (CLI) session with a TX Matrix Plus router in a routing matrix. In the example, the TX Matrix Plus router, which is also called the switch-fabric chassis (SFC), is known by the IP host name **host-sfc-0** and contains redundant Routing Engines. The commands display information about the management Ethernet interface and both internal Ethernet interfaces configured on the Routing Engine to which you are currently logged in:

```
user@host-sfc-0> show interfaces em0 terse
Interface               Admin Link Proto    Local                 Remote
em0                     up    up
em0.0                   up    up   inet     192.168.35.95/24

user@host-sfc-0> show interfaces ixgbe0 terse
Interface               Admin Link Proto    Local                 Remote
ixgbe0                  up    up
ixgbe0.0                up    up   inet     10.34.0.4/8
                                            162.0.0.4/2
                                   inet6    fe80::200:ff:fe22:4/64
                                            fec0::a:22:0:4/64
                                   tnp      0x22000004

user@host-sfc-0> show interfaces ixgbe1 terse
Interface               Admin Link Proto    Local                 Remote
ixgbe1                  up    up
ixgbe1.0                up    up   inet     10.34.0.4/8
                                            162.0.0.4/2
                                   inet6    fe80::200:1ff:fe22:4/64
                                            fec0::a:22:0:4/64
                                   tnp      0x22000004
```

The following example is a sequence of **show interfaces** commands issued in a CLI session with a T1600 router in a routing matrix. In the example, the T1600 router, which is also called the line-card chassis (LCC), is known by the IP host name **host-sfc-0-lcc-2** and contains redundant Routing Engines.

This T1600 router is connected to the routing matrix through a connection in the TXP-SIB-F13 in slot 2 of the SCC. The commands display information about the management Ethernet interface and both internal Ethernet interfaces configured on the Routing Engine to which you are currently logged in:

NOTE: In a routing matrix, the **show interfaces** command displays information about the current router only. If you are logged in to the TX Matrix Plus router, the **show interfaces** command output does not include information about any of the attached T1600 routers. To display interface information about a specific T1600 router in the routing matrix, you must first log in to that router.

The previous example shows a CLI session with the TX Matrix Plus router. To display interface information about the T1600 router known as **host-sfc-0-lcc-2**, first use the **request routing-engine login** command to log in to that LCC.

```
user@host-sfc-0> request routing-engine login lcc 2
--- JUNOS 9.6I built 2009-06-22 18:13:04 UTC
% cli
warning: This chassis is a Line Card Chassis (LCC) in a multichassis system.
warning: Use of interactive commands should be limited to debugging.
warning: Normal CLI access is provided by the Switch Fabric Chassis (SFC).
warning: Please logout and log into the SFC to use CLI.


user@host-sfc-0-lcc-2> show interfaces em0 terse
Interface               Admin Link Proto    Local                 Remote
em0                     up    up
em0.0                   up    up   inet     192.168.35.117/24

user@host-sfc-0-lcc-2> show interfaces bcm0 terse
Interface               Admin Link Proto    Local                 Remote
bcm0                    up    up
bcm0.0                  up    up   inet     10.1.0.5/8
                                            129.0.0.5/2
                                   inet6    fe80::201:ff:fe01:5/64
                                            fec0::a:1:0:5/64
                                   tnp      0x1000005

user@host-sfc-0-lcc-2> show interfaces em1 terse
Interface               Admin Link Proto    Local                 Remote
em1                     up    up
em1.0                   up    up   inet     10.1.0.5/8
                                            129.0.0.5/2
                                   inet6    fe80::201:1ff:fe01:5/64
                                            fec0::a:1:0:5/64
                                   tnp      0x1000005
```

Related Documentation

- *Understanding Internal Ethernet Interfaces*

CHAPTER 3

Configuring the Management Ethernet Interface

Management Ethernet Interface Overview

The router's management Ethernet interface, **fxp0** or **em0**, is an out-of-band management interface that needs to be configured only if you want to connect to the router through the management port on the front of the router. You can configure an IP address and prefix length for this interface, which you commonly do when you first install the Junos OS:

```
[edit]
user@host# set interfaces (fxp0 | em0) unit 0 family inet address/prefix-length
[edit]
user@host# show
interfaces {
    (fxp0 | em0) {
        unit 0 {
            family inet {
                address/prefix-length;
            }
        }
    }
}
```

To determine which management interface type is supported on a router, locate the router and Routing Engine combination in *Supported Routing Engines by Router* and note its management Ethernet interface type, either **em0** or **fxp0**.

Related Documentation

- Configuring a Consistent Management IP Address on page 26
- Configuring the MAC Address on the Management Ethernet Interface on page 27
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 16

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Consistent Management IP Address

On routers with multiple Routing Engines, each Routing Engine is configured with a separate IP address for the management Ethernet interface. To access the master Routing Engine, you must know which Routing Engine is active and use the appropriate IP address.

Optionally, for consistent access to the master Routing Engine, you can configure an additional IP address and use this address for the management interface regardless of which Routing Engine is active. This additional IP address is active only on the management Ethernet interface for the master Routing Engine. During switchover, the address moves to the new master Routing Engine.

NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is **fxp0**. For TX Matrix Plus routers and T1600 or T4000 routers configured in a routing matrix, the management Ethernet interface is **em0**.

NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the **fxp0** management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the **fxp0** management Ethernet interface so that the commands reference the **em0** management Ethernet interface instead.

To configure an additional IP address for the management Ethernet interface, include the **master-only** statement at the **[edit groups]** hierarchy level.

In the following example, IP address **10.17.40.131** is configured for both Routing Engines and includes a **master-only** statement. With this configuration, the **10.17.40.131** address is active only on the master Routing Engine. The address remains consistent regardless of which Routing Engine is active. IP address **10.17.40.132** is assigned to **fxp0** on **re0**, and address **10.17.40.133** is assigned to **fxp0** on **re1**.

```
[edit groups re0 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.132/25;
  }
}
[edit groups re1 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
```



```
        address 10.17.40.133/25;
      }
    }
```

This feature is available on all routers that include dual Routing Engines. On the TX Matrix router, this feature is applicable to the switch-card chassis (SCC) only.

Related Documentation

- Management Ethernet Interface Overview on page 25
- Configuring the MAC Address on the Management Ethernet Interface on page 27

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the MAC Address on the Management Ethernet Interface

By default, the router's management Ethernet interface uses as its MAC address the MAC address that is burned into the Ethernet card.

NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is **fxp0**. For TX Matrix Plus routers and T1600 routers configured in a routing matrix, and TX Matrix Plus routers with 3D SIBs, T1600 routers, and T4000 routers configured in a routing matrix, the management Ethernet interface is **em0**.

NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the **fxp0** management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the **fxp0** management Ethernet interface so that the commands reference the **em0** management Ethernet interface instead.

To display the MAC address used by the router's management Ethernet interface, enter the **show interface fxp0** or **show interface em0** operational mode command.

To change the management Ethernet interface's MAC address, include the **mac** statement at the **[edit interfaces fxp0]** or **[edit interfaces em0]** hierarchy level:

```
[edit interfaces (fxp0 | em0)]
mac mac-address;
```

Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** (for example, **0011.2233.4455**) or ***nn:nn:nn:nn:nn:nn*** (for example, **00:11:22:33:44:55**).

NOTE: If you integrate a standalone T640 router into a routing matrix, the PIC MAC addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, if you integrate a standalone T1600 router into a routing matrix, the PIC MAC addresses for the integrated T1600 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the configuration of a formerly standalone T1600 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

Related Documentation

- Management Ethernet Interface Overview on page 25
- Configuring a Consistent Management IP Address on page 26
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 16

- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 4

Enabling Passive Monitoring on Ethernet Interfaces

Passive Monitoring on Ethernet Interfaces Overview

The Monitoring Services I and Monitoring Services II PICs are designed to enable IP services. You can monitor IPv4 traffic if you have a Monitoring Services PIC installed in the router with the following PICs:

- 10-port Gigabit Ethernet PIC with SFPs
- 4-port Gigabit Ethernet PIC with SFPs
- 2-port Gigabit Ethernet PIC with SFPs
- 1-port 10-Gigabit Ethernet PIC

NOTE: The PICs in the preceding list support only IPv4.

NOTE: Starting with Junos OS Release 9.5, I2.0 based M120 routers and I3.0 based M320 routers with the PICs in the preceding list support passive monitoring. Other M Series and T Series routers with the PICs listed above started supporting passive monitoring before Junos OS Release 7.3. Support for 1-port 10-Gigabit Ethernet PIC with XENPAK on I2.0-based M120 routers and I3.0-based M320 routers was added in Junos OS Release 9.5.

- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (T640, T1600, and T4000 Core Routers) (supported on both WAN-PHY and LAN-PHY modes for both IPv4 and IPv6 addresses)

The following interfaces support passive monitoring on the I3.0-based MX 240, MX 480, and MX 960 routers, starting with Junos OS Release 8.5:

- Type 2 MX FPCs
- Type 3 MX FPCs
- Gigabit Ethernet Enhanced DPC with SFP (DPCE-R-40GE-SFP)
- 4-port 10-Gigabit Ethernet Enhanced DPCs with XFP (DPCE-R-4XGE-XFP)

The following interfaces support passive monitoring on the Trio-based MX 240, MX 480, and MX 960 routers:

- 10-Gigabit Ethernet MPC with SFP+
- 30-Gigabit Ethernet MPC
- 60-Gigabit Ethernet MPC

Passive monitoring is also supported on MX 80 routers with 10-Gigabit Ethernet MPC with SFP+ and 30-Gigabit Ethernet MPC interfaces.

Interfaces configured on the following FPCs and PIC support IPv6 passive monitoring on the T640, T1600, and T4000 routers:

- Enhanced Scaling FPC2
- Enhanced Scaling FPC3
- Enhanced Scaling FPC4
- Enhanced Scaling FPC4.1
- Enhanced II FPC1 (T640 and T1600 routers)
- Enhanced II FPC2 (T640 and T1600 routers)
- Enhanced II FPC3 (T640 and T1600 routers)
- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (supported on both WAN-PHY and LAN-PHY modes for both IPv4 and IPv6 addresses)
- Gigabit Ethernet PIC with SFP
- 10-Gigabit Ethernet PIC with XENPAK (T640 and T1600 routers)
- SONET/SDH OC192/STM64 PICs with XFP (T1600 and T4000 routers)
- SONET/SDH OC48c/STM16 PIC with SFP
- SONET/SDH OC12/STM4 (Multi-Rate) PIC with SFP (T1600 router)
- Type 1 SONET/SDH OC3/STM1 (Multi-Rate) PIC with SFP (T1600 router)

NOTE: **Unlike IPv4 passive monitoring, IPv6 passive monitoring is not supported on Monitoring Services PICs. You must configure port mirroring to forward the packets from the passive monitored ports to other interfaces.**

Release History Table

Release	Description
9.5	Starting with Junos OS Release 9.5, I2.0 based M120 routers and I3.0 based M320 routers with the PICs in the preceding list support passive monitoring.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling Passive Monitoring on Ethernet Interfaces

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you include the **passive-monitor-mode** statement in the configuration:

- Gigabit and Fast Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.
- Ethernet interfaces do not support the **stacked-vlan-tagging** statement for both IPv4 and IPv6 packets in passive monitor mode.

To enable packet flow monitoring on Ethernet interfaces:

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Include the **passive-monitor-mode** statement.

   ```
   [edit interfaces interface-name]
   user@host# set passive-monitor-mode
   ```

For IPv4 monitoring services interfaces, enable packet flow monitoring by including the **family** statement at the **[edit interfaces mo-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level, specifying the **inet** option:

1. In configuration mode, go to the **[edit interfaces mo-fpc/pic/port unit logical-unit-number]** hierarchy level.

   ```
   [edit]
   user@host# edit interfaces mo-fpc/pic/port unit logical-unit-number
   ```

2. Include the **passive-monitor-mode** statement.

   ```
   [edit interfaces mo-fpc/pic/port unit logical-unit-number]
   ```

```
user@host# set family inet
```

For conformity with the cflowd record structure, you must include the **receive-options-packets** and **receive-ttl-exceeded** statements at the **[edit interfaces mo-*fpc/pic/port* unit *logical-unit-number* family inet]** hierarchy level:

1. In configuration mode, go to the **[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]** hierarchy level.

   ```
   [edit]
   user@host# edit interfaces mo-fpc/pic/port unit logical-unit-number family inet
   ```

2. Include the **receive-options-packets** and **receive-ttl-exceeded** statements.

   ```
   [edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]
   user@host# set receive-options-packets
   user@host# set receive-ttl-exceeded
   ```

IPv6 passive monitoring is not supported on monitoring services PICs. A user must configure port mirroring to forward the packets from the passive monitored ports to other interfaces.

For information on FPCs and PICs that support IPv6 passive monitoring on the T640, T1600, and T4000 routers, see "Passive Monitoring on Ethernet Interfaces Overview" on page 29. Interfaces configured on these FPCs and PICs support IPv6 passive monitoring.

To configure port mirroring, include the **port-mirroring** statement at the **[edit forwarding-options]** hierarchy level.

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see *Configuring Multiservice Physical Interface Properties* and the *Junos OS Services Interfaces Library for Routing Devices*.

Related Documentation

- Passive Monitoring on Ethernet Interfaces Overview on page 29
- *Configuring Multiservice Physical Interface Properties*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 5

Configuring IEEE 802.1x Port-Based Network Access Control

IEEE 802.1x Port-Based Network Access Control Overview

MX Series routers support the IEEE 802.1x Port-Based Network Access Control (dot1x) protocol on Ethernet interfaces for validation of client and user credentials to prevent unauthorized access to a specified router port. Before authentication is complete, only 802.1x control packets are allowed and forwarded to the router control plane for processing. All other packets are dropped.

Authentication methods used must be 802.1x compliant. Authentication using RADIUS and Microsoft Active Directory servers is supported. The following user/client authentication methods are allowed:

- EAP-MD5 (RFC 3748)
- EAP-TTLS requires a server certificate (RFC 2716)
- EAP-TLS requires a client and server certificate
- PEAP requires only a server certificate

You can use both client and server certificates in all types of authentication except EAP-MD5.

NOTE: On the MX Series router, 802.1x can be enabled on bridged ports only and not on routed ports.

Dynamic changes to a user session are supported to allow the router administrator to terminate an already authenticated session by using the “RADIUS disconnect” message defined in RFC 3576.

Related Documentation

- Understanding the Administrative State of the Authenticator Port on page 34
- Understanding the Administrative Mode of the Authenticator Port on page 34
- Configuring the Authenticator on page 35
- Viewing the dot1x Configuration on page 35
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding the Administrative State of the Authenticator Port

The administrative state of an authenticator port can take any of the following three states:

- Force authorized—Allows network access to all users of the port without requiring them to be authenticated. This is equivalent to not having any authentication enabled on the port.
- Force unauthorized—Denies network access to all users of the port. This is equivalent to disabling the port.
- Automatic—This is the default mode where the authentication server response determines if the port is opened for traffic or not. Only the successfully authenticated clients are allowed access, all others are denied.

In Junos OS, the default mode is “automatic.” The “force authorized” and “force unauthorized” admin modes are not supported. You can achieve the functionality of “force authorized” mode by disabling **dot1x** on the required port. You can achieve the functionality of “force unauthorized” mode by disabling the port itself.

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- Understanding the Administrative Mode of the Authenticator Port on page 34
- Configuring the Authenticator on page 35
- Viewing the dot1x Configuration on page 35
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding the Administrative Mode of the Authenticator Port

Junos OS supports the supplicant mode “single” and not the “single secure” nor “multiple” modes. The “Single” mode option authenticates only the first client that connects to a port. All other clients that connect later (802.1x compliant or noncompliant) are allowed free access on that port without any further authentication. If the first authenticated client logs out, all other users are locked out until a client authenticates again.

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- Understanding the Administrative State of the Authenticator Port on page 34
- Configuring the Authenticator on page 35
- Viewing the dot1x Configuration on page 35
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the Authenticator

To configure the IEEE 802.1x Port-Based Network Access Control protocol on Ethernet interfaces you must configure the **authenticator** statement at the **[edit protocols dot1x]** hierarchy level. Use the **authentication-profile-name** ***access-profile-name*** statement to specify the authenticating RADIUS server, and use the **interface** statement to specify and configure the Gigabit Ethernet or Fast Ethernet interface on the router specifically for IEEE 802.1x protocol use; both at the **[edit protocols dot1x authenticator]** hierarchy level.

```
[edit protocols dot1x]
authenticator {
  authentication-profile-name access-profile-name;
  interface (xe-fpc/pic/port | ge-fpc/pic/port | fe-fpc/pic/port) {
    maximum-requests seconds;
    quiet-period seconds;
    reauthentication (disable | interval seconds);
    retries integer;
    server-timeout seconds;
    supplicant (single);
    supplicant-timeout seconds;
    transmit-period seconds;
  }
}
```

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- Understanding the Administrative State of the Authenticator Port on page 34
- Understanding the Administrative Mode of the Authenticator Port on page 34
- Viewing the dot1x Configuration on page 35
- *Ethernet Interfaces Feature Guide for Routing Devices*

Viewing the dot1x Configuration

Purpose To review and verify the dot1x configuration.

Action To view all **dot1x** configurations, use the **show dot1x interface** operational mode command. To view a **dot1x** configuration for a specific interface, use the **show dot1x interface (xe-*fpc/pic/port* | ge-*fpc/pic/port* | fe-*fpc/pic/port*) detail** operational mode command.

See the *Network Interfaces Command Reference* for more information about this command.

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- Understanding the Administrative State of the Authenticator Port on page 34
- Understanding the Administrative Mode of the Authenticator Port on page 34
- Configuring the Authenticator on page 35
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 6

Configuring IEEE 802.1x Port-Based Network Access Control in Enhanced LAN Mode

802.1X for MX Series Routers in Enhanced LAN Mode Overview

Starting with Junos os Release 14.2, IEEE 802.1X provides network edge security, protecting Ethernet LANs from unauthorized user access. Support is implemented for controlling access to your network through an MX Series router by using several different authentication methods, such as 802.1X, MAC RADIUS, or a captive portal.

This functionality is supported on the following MPCs on MX240, MX480, and MX960 routers in enhanced LAN mode:

- MPC4E with two 100-Gigabit Ethernet ports and eight 10-Gigabit Ethernet ports
- MPC4E with thirty-two 10-Gigabit Ethernet ports
- MPC3E that contains a 2-port 40-Gigabit Ethernet MIC with QSFP+
- MPC1E with forty 1-Gigabit Ethernet ports or twenty 1-Gigabit Ethernet ports

You must reboot the router when you configure or delete the enhanced LAN mode on the router. Configuring the **network-services lan** option implies that the system is running in the enhanced IP mode. When you configure a device to function in MX-LAN mode, only the supported configuration statements and operational show commands that are available for enabling or viewing in this mode are displayed in the CLI interface. If your system contains parameters that are not supported in MX-LAN mode in a configuration file, you cannot commit those unsupported attributes. You must remove the settings that are not supported and then commit the configuration. After the successful CLI commit, a system reboot is required for the attributes to be come effective. Similarly, if you remove the **network-services lan** statement, the system does not run in MX-LAN mode. Therefore, all of the settings that are supported outside of the MX-LAN mode are displayed and are available for definition in the CLI interface. If your configuration file contains settings that are supported only in MX-LAN mode, you must remove those attributes before you commit the configuration. After the successful CLI commit, a system reboot will be required for the CLI settings to take effect. The Layer 2 Next-Generation CLI configuration settings are supported in MX-LAN mode. As a result, the typical MX Series-format of CLI configurations might differ in MX-LAN mode.

This functionality is supported on an MX Series Virtual Chassis combination that functions in enhanced LAN mode (by entering the **network-services lan** statement at the **[edit chassis]** hierarchy level). Port-based network access control is supported on MX240, MX480, and MX960 routers with MPCs in both the MX-LAN mode and the non-MX-LAN mode (with other supported network services modes on MPCs on these routers). To configure the IEEE 802.1x port-based network access control (PNAC) protocol on Ethernet interfaces, you must configure the **authenticator** statement at the **[edit protocols authentication-access- control]** hierarchy level. You can also configure captive portal authentication on a router so that users connected to the switch are authenticated before being allowed to access the network. You can also configure Junos Pulse Access Control Service as the access policy to authenticate and authorize users connected to the switch for admission to the network and for access to protected network resources by using the **uac-policy** statement.

How 802.1X Authentication Works

802.1X authentication works by using an *Authenticator Port Access Entity* (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the *Authentication server* (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant.

The end device is authenticated in either *single* mode, *single-secure* mode, or *multiple* mode:

- **single**—Authenticates only the first end device. All other end devices that connect later to the port are allowed full access without any further authentication. They effectively "piggyback" on the end devices' authentication.
- **single-secure**—Allows only one end device to connect to the port. No other end device is allowed to connect until the first logs out.
- **multiple**—Allows multiple end devices to connect to the port. Each end device will be authenticated individually.

Network access can be further defined using VLANs and firewall filters, which both act as filters to separate and match groups of end devices to the areas of the LAN they require. For example, you can configure VLANs to handle different categories of authentication failures depending upon:

- Whether or not the end device is 802.1X-enabled.
- Whether or not MAC RADIUS authentication has been configured on the switch interfaces to which the hosts are connected.
- Whether the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message. See *Configuring RADIUS Server Fail Fallback (CLI Procedure)*.

802.1X Features Overview

NOTE: The 802.1X features available on the MX Series routers depend upon which switch you are using.

802.1X features on Juniper Networks MX Series routers are:

- Guest VLAN—Provides limited access to a LAN, typically just to the Internet, for nonresponsive end devices that are not 802.1X-enabled when MAC RADIUS authentication has not been configured on the switch interfaces to which the hosts are connected . Also, a guest VLAN can be used to provide limited access to a LAN for guest users. Typically, the guest VLAN provides access just to the Internet and to other guests' end devices.
- Server-reject VLAN—Provides limited access to a LAN, typically just to the Internet, for responsive end devices that are 802.1X-enabled but that have sent the wrong credentials.
- Server-fail VLAN—Provides limited access to a LAN, typically just to the Internet, for 802.1X end devices during a RADIUS server timeout.

- Dynamic VLAN—Enables an end device, after authentication, to be a member of a VLAN dynamically.
- Private VLAN—Enables configuration of 802.1X authentication on interfaces that are members of private VLANs (PVLANs).
- Dynamic changes to a user session—Allows the switch administrator to terminate an already authenticated session. This feature is based on support of the RADIUS Disconnect Message defined in RFC 3576.
- RADIUS accounting—Sends accounting information to the RADIUS accounting server. Accounting information is sent to the server whenever a subscriber logs in or logs out and whenever a subscriber activates or deactivates a subscription.

Supported Features Related to 802.1X Authentication

802.1X does not replace other security technologies. 802.1X works together with port security features, such as DHCP snooping, dynamic ARP inspection (DAI), and MAC limiting, to guard against spoofing.

Supported features related to authentication include:

- Static MAC bypass—Provides a bypass mechanism to authenticate devices that are not 802.1X-enabled (such as printers). Static MAC bypass connects these devices to 802.1X-enabled ports, bypassing 802.1X authentication.
- MAC RADIUS authentication—Provides a means to enable or disable MAC authentication independently of whether 802.1X authentication is enabled.

Release History Table

Release	Description
14.2	Starting with Junos os Release 14.2, IEEE 802.1X provides network edge security, protecting Ethernet LANs from unauthorized user access. Support is implemented for controlling access to your network through an MX Series router by using several different authentication methods, such as 802.1X, MAC RADIUS, or a captive portal.

Understanding 802.1X and LLDP and LLDP-MED on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, Juniper Networks MX Series routers use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the router to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include information such as chassis and port identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in the Juniper Networks Junos operating system (Junos OS).

LLDP-MED goes one step further than LLDP, exchanging IP-telephony messages between the router and the IP telephone.

LLDP and LLDP-MED also provide PoE power management capabilities. LLDP power negotiation allows the router to manage PoE power by negotiating with LLDP-enabled powered devices to dynamically allocate PoE power as needed. LLDP power priority allows an LLDP-enabled powered device to set the PoE power priority on the router interface to which it connects.

The router also uses these protocols to ensure that voice traffic gets tagged and prioritized with the correct values at the source itself. For example, 802.1p CoS and 802.1Q tag information can be sent to the IP telephone.

EX Series routeres support the following basic TLVs:

- **Chassis Identifier**—The MAC address associated with the local system.

  NOTE: **The Chassis ID TLV has a subtype for Network Address Family. LLDP frames are validated only if this subtype has a value of 1 (IPv4) or 2 (IPv6). For any other value, the transmitting device is detected by LLDP as a neighbor and displayed in the output of the "show lldp neighbors" command, but is not assigned to the VLAN.**

- **Port Identifier**—The port identification for the specified port in the local system.
- **Port Description**—Textual description of the interface or the logical unit. The description for the logical unit is used, if available; otherwise, the Port Description TLV will contain the description configured on the physical interface. For example, LAG member interfaces do not contain a logical unit, so only the description configured on the physical interface can be used.
- **System Name**—The user-configured name of the local system. The system name can be a maximum of 256 characters.
- **System Description**—The system description containing information about the software and current image running on the system. This information is not configurable, but taken from the software.
- **System Capabilities**—The primary function performed by the system. The capabilities that system supports; for example, bridge or router. This information is not configurable, but based on the model of the product.
- **Management Address**—The IPv4 or IPv6 management address of the local system.

EX Series routeres support the following 802.3 TLVs:

- **Power via MDI**—A TLV that advertises MDI power support, PSE power pair, and power class information.
- **MAC/PHY Configuration Status**—A TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is not configurable, but based on the physical interface structure.

NOTE: The MAC/PHY Configuration Status TLV has a subtype for the PMD Auto-Negotiation Advertised Capability field. This field will contain a value of other or unknown if the LLDP packet was transmitted from a 10-gigabit SFP+ port.

- **Link Aggregation**—A TLV that advertises if the port is aggregated and its aggregated port ID.
- **Maximum Frame Size**—A TLV that advertises the Maximum Transmission Unit (MTU) of the interface sending LLDP frames.
- **Port Vlan**—A TLV that advertises the VLAN name configured on the interface.

EX Series routeres support the following LLDP-MED TLVs:

- **LLDP MED Capabilities**—A TLV that advertises the primary function of the port. The capabilities values range 0 through 15:
  - **0**— Capabilities
  - **1**— Network Policy
  - **2**— Location Identification
  - **3**— Extended Power via MDI-PSE
  - **4**— Inventory
  - **5–15**— Reserved
- LLDP-MED Device Class Values:
  - **0**— Class not defined.
  - **1**— Class 1 Device.
  - **2**— Class 2 Device.
  - **3**— Class 3 Device.
  - **4**— Network Connectivity Device
  - **5–255**— Reserved.
- **Network Policy**—A TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types, such as voice or streaming video, 802.1Q VLAN tagging, and 802.1p priority bits and Diffserv code points.
- **Endpoint Location**— A TLV that advertises the physical location of the endpoint.
- **Extended Power via MDI**— A TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, Juniper Networks MX Series routers use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the router to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently.

Understanding 802.1X and RADIUS Accounting on MX Series Routers in Enhanced LAN Mode

Juniper Networks MX Series routers support IETF RFC 2866, *RADIUS Accounting*. Starting with Junos OS Release 14.2, you can configure RADIUS accounting on an MX Series router which enables statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

To configure RADIUS accounting, specify one or more RADIUS accounting servers to receive the statistical data from the switch, and select the type of accounting data to be collected.

The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify a list of RADIUS accounting servers. In the event that the primary server (the first one configured) is unavailable, each RADIUS server in the list is tried in the order in which they are configured in the Juniper Networks Junos operating system (Junos OS).

The RADIUS accounting process between a switch and a RADIUS server works like this:

1. A RADIUS accounting server listens for User Datagram Protocol (UDP) packets on a specific port. For example, on FreeRADIUS, the default port is 1813.

2. The switch forwards an *accounting-request* packet containing an event record to the accounting server. For example, a supplicant is authenticated through 802.1X authentication and connected to the LAN. The event record associated with this supplicant contains an *Acct-Status-Type* attribute whose value indicates the beginning of user service for this supplicant. When the supplicant's session ends, the accounting request will contain an *Acct-Status-Type* attribute value indicating the end of user service. The RADIUS accounting server records this as a stop-accounting record containing session information and the length of the session.

3. The RADIUS accounting server logs these events as start-accounting or stop-accounting records. The records are in a file. On FreeRADIUS, the file name is the server's address; for example, 122.69.1.250.

4. The accounting server sends an *accounting-response* packet back to the switch confirming it has received the accounting request.

5. If the switch does not receive a response from the server, it continues to send accounting requests until an accounting response is returned from the accounting server.

The statistics collected through this process can be displayed from the RADIUS server; to see those statistics, the user accesses the log file configured to receive them.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can configure RADIUS accounting on an MX Series router which enables statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

Understanding 802.1X and VoIP on MX Series Routers in Enhanced LAN Mode

When you use Voice over IP (VoIP), you can connect IP telephones to the router and configure IEEE 802.1X authentication for 802.1X-compatible IP telephones. Starting with Junos OS Release 14.2, 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access.

VoIP is a protocol used for the transmission of voice through packet-switched networks. VoIP transmits voice calls using a network connection instead of an analog phone line.

When VoIP is used with 802.1X, the RADIUS server authenticates the phone, and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) provides the class-of-service (CoS) parameters to the phone.

You can configure 802.1X authentication to work with VoIP in multiple supplicant or single supplicant mode. In *multiple-supplicant* mode, the 802.1X process allows multiple supplicants to connect to the interface. Each supplicant will be authenticated individually. For an example of a VoIP multiple supplicant topology, see Figure 1 on page 46.

Figure 1: VoIP Multiple Supplicant Topology

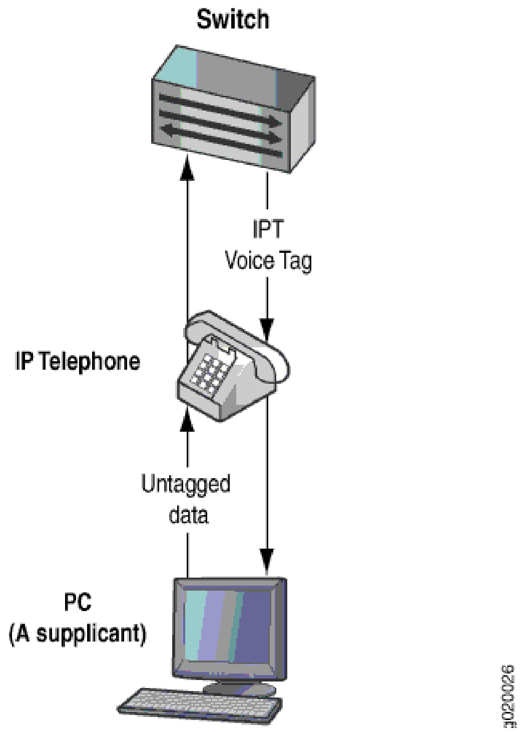

If an 802.1X-compatible IP telephone does not have an 802.1X host but has another 802.1X-compatible device connected to its data port, you can connect the phone to an interface in single-supplicant mode. In *single-supplicant* mode, the 802.1X process authenticates only the first supplicant. All other supplicants who connect later to the interface are allowed full access without any further authentication. They effectively “piggyback” on the first supplicant’s authentication. For an example of a VoIP single supplicant topology, see Figure 2 on page 47 .

Figure 2: VoIP Single Supplicant Topology

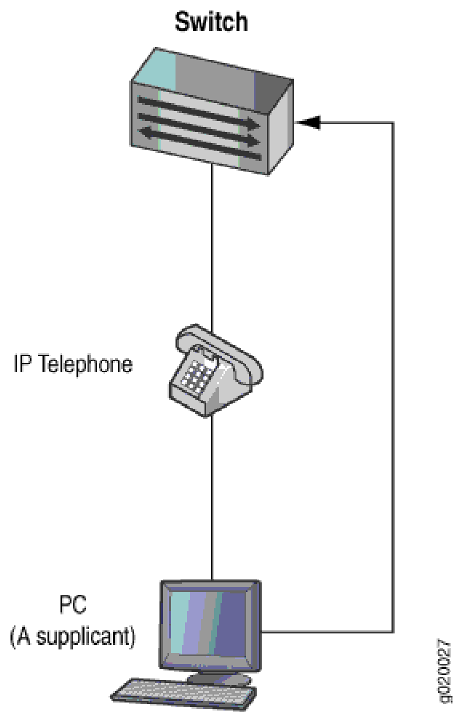

If an IP telephone does not support 802.1X, you can configure VoIP to bypass 802.1X and LLDP-MED and have the packets forwarded to a VoIP VLAN,

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access.

Understanding Guest VLANs for 802.1X on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, guest VLANs can be configured on switches that are using 802.1X authentication to provide limited access—typically only to the Internet—for:

- Corporate guests
- End devices that are not 802.1X-enabled
- Nonresponsive end devices when MAC RADIUS authentication has not been configured on the switch interfaces to which the hosts are connected

A guest VLAN is not used for supplicants sending incorrect credentials. Those supplicants are directed to the server-reject VLAN instead.

For end devices that are not 802.1X-enabled, a guest VLAN can allow limited access to a server from which the non-802.1X-enabled end device can download the supplicant software and attempt authentication again.

A guest VLAN is not used when MAC RADIUS authentication has been configured on the switch interfaces to which the hosts are connected. Some end devices, such as a printer, cannot be enabled for 802.1X. The hosts for such devices should be connected to switch interfaces that are configured for MAC RADIUS authentication.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, guest VLANs can be configured on switches that are using 802.1X authentication to provide limited access—typically only to the Internet

Understanding Dynamic VLANs for 802.1X on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, dynamic VLANs, in conjunction with the 802.1X authentication process, provide secure access to the LAN for end devices belonging to different VLANs on a single port.

When this feature is configured on the RADIUS server, an end device or user authenticating on the RADIUS server is assigned to the VLAN configured for it. The end device or user becomes a member of a VLAN dynamically after successful 802.1X authentication. For information on configuring dynamic VLANs on your RADIUS server, see the documentation for your RADIUS server.

Successful authentication requires that the VLAN ID or VLAN name exist on the router and match the VLAN ID or VLAN name sent by the RADIUS server during authentication. If neither exists, the end device is unauthenticated. If a guest VLAN is established, the unauthenticated end device is automatically moved to the guest VLAN.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, dynamic VLANs, in conjunction with the 802.1X authentication process, provide secure access to the LAN for end devices belonging to different VLANs on a single port.

Understanding Server Fail Fallback and Authentication on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, server fail fallback allows you to specify how end devices connected to the router are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

Juniper Networks MX Series routers in enhanced LAN mode use authentication to implement access control in an enterprise network. If 802.1X, MAC RADIUS, or captive portal authentication are configured on the interface, end devices are evaluated at the initial connection by an authentication (RADIUS) server. If the end device is configured on the authentication server, the device is granted access to the LAN and the MX Series router opens the interface to permit access.

A RADIUS server timeout occurs if no RADIUS authentication servers are reachable when an end device logs in and attempts to access the LAN. Server fail fallback allows you to specify one of four actions to be taken toward end devices awaiting authentication when the server is timed out:

- *Permit* authentication, allowing traffic to flow from the end device through the interface as if the end device were successfully authenticated by the RADIUS server.
- *Deny* authentication, preventing traffic from flowing from the end device through the interface. This is the default.
- *Move* the end device to a specified VLAN. (The VLAN must already exist on the router.)
- *Sustain* authenticated end devices that already have LAN access and *deny* unauthenticated end devices. If the RADIUS servers time out during reauthentication, previously authenticated end devices are reauthenticated and new users are denied LAN access.

Server fail fallback is triggered most often during reauthentication when the already configured and in-use RADIUS server becomes inaccessible. However, server fail fallback can also be triggered by an end device's first attempt at authentication through the RADIUS server.

Server fail fallback allows you to specify that an end device be moved to a specified VLAN if the router receives a RADIUS access-reject message. The configured VLAN name overrides any attributes sent by the server.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, server fail fallback allows you to specify how end devices connected to the router are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

Configuring 802.1X RADIUS Accounting on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, RADIUS accounting permits statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

To configure basic RADIUS accounting using the CLI:

1. Specify the accounting servers to which the switch will forward accounting statistics:

   ```
   [edit access ]
   user@router# set profile profile1 radius accounting-serveraccounting-server [122.69.1.250
   122.69.1.252]
   ```

2. Define the RADIUS accounting servers:

   ```
   [edit access]
   user@router# set radius-server 122.69.1.250 secret juniper
   user@router# set radius-server 122.69.1.252 secret juniper1
   ```

3. Enable accounting for an access profile:

   ```
   [edit access]
   user@router# set profile profile1 accounting
   ```

4. Configure the RADIUS servers to use while sending accounting messages and updates:

   ```
   [edit access]
   user@router# set profile profile1 accounting order radius
   ```

5. Configure the statistics to be collected on the router and forwarded to the accounting server:

   ```
   [edit access ]
   user@router# set profile profile1 accounting accounting-stop-on-access-deny
   user@router# set profile profile1 accounting accounting-stop-on-failure
   ```

6. Display accounting statistics collected on the router:

   ```
   user@router> show network-access aaa statistics accounting
   Accounting module statistics
     Requests received: 1
     Accounting Response failures: 0
   ```



```
Accounting Response Success: 1
Requests timedout: 0
```

7. Open an accounting log on the RADIUS accounting server using the server's address, and view accounting statistics:

```
[root@freeradius]# cd /usr/local/var/log/radius/radacct/122.69.1.250
[root@freeradius 122.69.1.250]# ls

detail-20071214


[root@freeradius 122.69.1.250]# vi details-20071214

        User-Name = "000347e1bab9"
        NAS-Port = 67
        Acct-Status-Type = Stop
        Acct-Session-Id = "802.1x811912"
        Acct-Input-Octets = 17454
        Acct-Output-Octets = 4245
        Acct-Session-Time = 1221041249
        Acct-Input-Packets = 72
        Acct-Output-Packets = 53
        Acct-Terminate-Cause = Lost-Carrier
        Acct-Input-Gigawords = 0
        Acct-Output-Gigawords = 0
        Called-Station-Id = "00-19-e2-50-52-60"
        Calling-Station-Id = "00-03-47-e1-ba-b9"
        Event-Timestamp = "Sep 10 2008 16:52:39 PDT"
        NAS-Identifier = "esp48t-1b-01"
        NAS-Port-Type = Virtual

        User-Name = "000347e1bab9"
        NAS-Port = 67
        Acct-Status-Type = Start
        Acct-Session-Id = "802.1x811219"
        Called-Station-Id = "00-19-e2-50-52-60"
        Calling-Station-Id = "00-03-47-e1-ba-b9"
        Event-Timestamp = "Sep 10 2008 18:58:52 PDT"
        NAS-Identifier = "esp48t-1b-01"
        NAS-Port-Type = Virtual
```

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, RADIUS accounting permits statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server.

Related Documentation

Configuring 802.1X Interface Settings on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, IEEE 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from a supplicant (client) at the interface until the supplicant's credentials are presented and matched on the *authentication server* (a RADIUS server). When the supplicant is authenticated, the switch stops blocking access and opens the interface to the supplicant.

NOTE:

- You can also specify an 802.1X exclusion list to specify supplicants can that can bypass authentication and be automatically connected to the LAN.
- You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.
- You cannot configure 802.1X user authentication on redundant trunk groups (RTGs).

Before you begin, specify the RADIUS server or servers to be used as the authentication server.

To configure 802.1X on an interface:

1. Configure the supplicant mode as **single** (authenticates the first supplicant), **single-secure** (authenticates only one supplicant), or **multiple** (authenticates multiple supplicants):

   ```
   [edit protocols authentication-access-control]
   user@switch# set interface ge-0/0/5 supplicant multiple
   ```

2. Enable reauthentication and specify the reauthentication interval:

   ```
   [edit protocols authentication-access-control]
   user@switch# set interface ge-0/0/5/0 dot1x reauthentication interval 5
   ```

3. Configure the interface timeout value for the response from the supplicant:

   ```
   [edit protocols authentication-access-control]
   user@switch# set interface ge-0/0/5 dot1x supplicant-timeout 5
   ```

4. Configure the timeout for the interface before it resends an authentication request to the RADIUS server:

   ```
   [edit protocols authentication-access-control]
   user@switch# set interface ge-0/0/5 server-timeout 5
   ```

5. Configure how long, in seconds, the interface waits before retransmitting the initial EAPOL PDUs to the supplicant:


```
[edit protocols authentication-access-control]
user@switch# set interface ge-0/0/5 dot1x transmit-period 60
```

6. Configure the maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out:

```
[edit protocols authentication-access-control]
user@switch# set interface ge-0/0/5 dot1x maximum-requests 5
```

7. Configure the number of times the switch attempts to authenticate the port after an initial failure. The port remains in a wait state during the quiet period after the authentication attempt.

```
[edit protocols authentication-access-control]
user@switch# set interface ge-0/0/5 retries 1
```

NOTE: This setting specifies the number of tries before the switch puts the interface in a "HELD" state.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, IEEE 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from a supplicant (client) at the interface until the supplicant's credentials are presented and matched on the *authentication server* (a RADIUS server).

Related Documentation

Configuring LLDP-MED on MX Series Routers in Enhanced LAN Mode

Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) is an extension of LLDP. Starting with Junos OS Release 14.2, the router uses LLDP-MED to support device discovery of VoIP telephones and to create location databases for these telephone locations.

LLDP-MED is turned on by default on MX Series routers.

This topic describes:

- Enabling LLDP-MED on Interfaces on page 54
- Configuring Location Information Advertised by the Router on page 54
- Configuring for Fast Start on page 54

Enabling LLDP-MED on Interfaces

LLDP-MED is enabled on all interfaces by default. If it is disabled, you can enable LLDP-MED by configuring it on all interfaces or on specific interfaces.

To configure LLDP-MED on all interfaces or on a specific interface:

```
[edit protocols lldp-med]
user@router# set interface (LLDP-MED) ge-0/0/2.0
```

Configuring Location Information Advertised by the Router

You can configure the location information that is advertised from the router to the LLDP-MED device. You can specify a civic-based location (geographic location) or a location based on an ELIN (Emergency Location Identification Number):

- To specify a location by geography:

```
[edit protocols lldp-med]
user@router# set interface ge-0/0/2.0 location civic-based country-code US
user@router# set interface ge-0/0/2.0 location civic-based ca-type 1 ca-value "El Dorado County"
user@router# set interface ge-0/0/2.0 location civic-based ca-type 2 ca-value CA
user@router# set interface ge-0/0/2.0 location civic-based ca-type 3 ca-value Somerset
user@router# set interface ge-0/0/2.0 location civic-based ca-type 6 ca-value "Mount Aukum Road"
user@router# set interface ge-0/0/2.0 location civic-based ca-type 19 ca-value 6450
user@router# set interface ge-0/0/2.0 location civic-based ca-type 21 ca-value "Holiday Market"
```

- To specify a location using an **elin** string:

```
[edit protocols lldp-med]
user@router# set interface ge-0/0/2.0 location elin 4085551212
```

Configuring for Fast Start

You can specify the number of LLDP-MED advertisements sent from the router in the first second after it has detected an LLDP-MED device. The default is 3; to set it to another value:

```
[edit protocols lldp-med]
user@router# set fast-start 6
```

NOTE: If an interface is configured as a VoIP interface, then the router does not wait for an attached phone to identify itself as an LLDP-MED device before it performs an LLDP-MED fast start after a graceful Routing Engine switchover (GRES) or a reboot. Instead, it immediately performs an LLDP-MED fast start after a GRES or reboot. This behavior prevents certain models of IP phones from resetting after a GRES.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, the router uses LLDP-MED to support device discovery of VoIP telephones and to create location databases for these telephone locations.

Related Documentation

Configuring LLDP on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, devices use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links.The information enables the device to quickly identify a variety of other devices, resulting in a LAN that interoperates smoothly and efficiently.

This topic describes:

- Enabling LLDP on Interfaces on page 55
- Adjusting LLDP Advertisement Settings on page 56
- Adjusting SNMP Notification Settings of LLDP Changes on page 57
- Specifying a Management Address for the LLDP Management TLV on page 57

Enabling LLDP on Interfaces

LLDP is enabled on all interfaces by default. If it is disabled, you can enable LLDP by configuring it on all interfaces or on specific interfaces.

- To configure LLDP on all interfaces:

```
[edit protocols lldp]
user@router# set interface all
```

- To configure LLDP on a specific interface:

```
[edit protocols lldp]
user@router# set interface interface-name
```

NOTE: On MX Series routers, LLDP cannot be configured on the management Ethernet interface. Issuing the command **set protocols lldp interfaceem0** generates the following error message:

```
error: name: 'em0': Invalid interface
error: statement creation failed: interface
```

Adjusting LLDP Advertisement Settings

You can adjust the following settings for LLDP advertisements for troubleshooting or verification purposes. The default values are applied when LLDP is enabled. For normal operations, we recommend that you do not change the default values.

- To specify the frequency at which LLDP advertisements are sent (in seconds):

  ```
  [edit protocols lldp]
  user@router# set advertisement-interval seconds
  ```

  For example, using the default value:

  ```
  [edit protocols lldp]
  user@router# set advertisement-interval 45
  ```

- To specify the number of seconds that LLDP information is held before it is discarded (the multiplier value is used in combination with the **advertisement-interval** value):

  ```
  [edit protocols lldp]
  user@router# set hold-multiplier seconds
  ```

  For example, using the default value:

  ```
  [edit protocols lldp]
  user@router# set hold-multiplier 5
  ```

- To specify the number of seconds the device delays before sending advertisements to neighbors after a change is made in a TLV (type, length, or value) element in LLDP or in the state of the local system, such as a change in hostname or management address, set the transmit delay. The transmit delay is enabled by default on switches to reduce the delay in notifying neighbors of a change in the local system. The default value is 2 seconds (if the **advertisement-interval** value is set to 8 seconds or more) or 1 second (if the **advertisement-interval** value is set to less than 8 seconds).

  ```
  [edit protocols lldp]
  user@router# set transmit-delay seconds
  ```

  For example:

  ```
  [edit protocols lldp]
  user@router# set transmit-delay 2
  ```

NOTE: The **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value; otherwise, an error is returned when you attempt to commit the configuration.

Adjusting SNMP Notification Settings of LLDP Changes

You can adjust the following settings for SNMP notifications of LLDP changes. If the values are not specified or if the interval values are set to **0**, the notifications are disabled.

- To specify the frequency at which LLDP database changes are sent (in seconds):

```
[edit protocols lldp]
user@router# set lldp-configuration-notification-interval seconds
```

  For example:

```
[edit protocols lldp]
user@router# set lldp-configuration-notification-interval 600
```

- To configure the time interval for SNMP trap notifications to wait for topology changes (in seconds):

```
[edit protocols lldp]
user@router# set ptopo-configuration-trap-interval seconds
```

  For example:

```
[edit protocols lldp]
user@router# set ptopo-configuration-trap-interval 600
```

- To specify the holding time (used in combination with the **ptopo-configuration-trap-interval** value) to maintain dynamic topology entries (in seconds):

```
[edit protocols lldp]
user@router# set ptopo-configuration-maximum-hold-time seconds
```

  For example:

```
[edit protocols lldp]
user@router# set ptopo-configuration-maximum-hold-time 2147483647
```

Specifying a Management Address for the LLDP Management TLV

You can configure an IPv4 or IPv6 management address to be used in the LLDP Management Address type, length, and value (TLV) messages. Only out-of-band management addresses must be used as the value for the **management-address** statement.

To configure the management address:

```
[edit protocols lldp]
user@router# set management-address ip-address
```

NOTE: Ensure that the interface with the configured management address has LLDP enabled using the **set protocols lldp interface** command. If you configure a customized management address for LLDP on an interface that has LLDP disabled, the **show lldp local-information** command output will not display the correct interface information.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, devices use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links.

Configuring Server Fail Fallback on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, server fail fallback allows you to specify how end devices connected to the router are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

802.1X and MAC RADIUS authentication work by using an *authenticator port access entity* (the router) to block all traffic to and from an end device at the interface until the end device's credentials are presented and matched on the *authentication server* (a RADIUS server). When the end device has been authenticated, the router stops blocking and opens the interface to the end device.

When you set up 802.1X or MAC RADIUS authentication on the router, you specify a primary authentication server and one or more backup authentication servers. If the primary authentication server cannot be reached by the router and the secondary authentication servers are also unreachable, a RADIUS server timeout occurs. Because the authentication server grants or denies access to the end devices awaiting authentication, the router does not receive access instructions for end devices attempting access to the LAN and normal authentication cannot be completed. Server fail fallback allows you to configure authentication alternatives that permit the router to take appropriate actions toward end devices awaiting authentication or reauthentication.

NOTE: The authentication fallback method called *server-reject VLAN* provides limited access to a LAN, typically just to the Internet, for responsive end devices that are 802.1X-enabled but that have sent the wrong credentials. If the end device that is authenticated using the server-reject VLAN is an IP phone, voice traffic is not allowed.

To configure basic server fail fallback options using the CLI:

- Configure an interface to allow traffic to flow from a supplicant to the LAN if a RADIUS server timeout occurs (as if the end device had been successfully authenticated by a RADIUS server):

```
[edit protocols authentication-access-control]
user@router# set interface ge-0/0/1 dot1x server-fail permit
```

- Configure an interface to prevent traffic flow from an end device to the LAN (as if the end device had failed authentication and had been rejected by the RADIUS server):

```
[edit protocols authentication-access-control]
user@router# set interface ge-0/0/1 dot1x server-fail deny
```

- Configure an interface to move an end device to a specified VLAN if a RADIUS server timeout occurs (in this case, the VLAN name is **vlan1**):

```
[edit protocols authentication-access-control]
user@router# set interface ge-0/0/1 dot1x server-fail vlan-name vlan1
```

- Configure an interface to recognize already connected end devices as reauthenticated if there is a RADIUS timeout during reauthentication (new users will be denied access):

```
[edit protocols authentication-access-control]
user@router# set interface ge-0/0/1 dot1x server-fail use-cache
```

- Configure an interface that receives a RADIUS access-reject message from the authentication server to move end devices attempting LAN access on the interface to a specified VLAN already configured on the router (in this case, the VLAN name is **vlan-sf**):

```
[edit protocols authentication-access-control]
user@router# set interface ge-0/0/1 dot1x server-reject-vlan vlan-sf
```

NOTE: If an IP phone is authenticated in the server-reject VLAN, voice traffic is not allowed.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, server fail fallback allows you to specify how end devices connected to the router are supported if the RADIUS authentication server becomes unavailable or sends a RADIUS access-reject message.

Understanding Captive Portal Authentication on the MX Series Routers

Starting with Junos OS Release 14.2, captive portal authentication (hereafter referred to as captive portal) allows you to authenticate users on MX Series routers by redirecting Web browser requests to a login page that requires users to input a username and password before they are allowed access to the network.Captive portal controls network access by requiring users to provide information that is authenticated against a RADIUS server database using EAP-MD5, You can also use captive portal to display an acceptable-use policy to users before they access your network.

Juniper Networks Junos Software for MX Series routers provides a template that allows you to easily design and modify the look of the captive portal login page. You enable specific interfaces for captive portal. The first time a client connected to a captive portal interface attempts to access a webpage, the switch presents the captive portal login page. Upon successful authentication, the user is allowed access to the network and to continue to the original page requested.

NOTE: **If Hypertext Transfer Protocol Secure (HTTPS) is enabled, Hypertext Transfer Protocol (HTTP) requests are redirected to an HTTPS connection for the captive portal authentication process. After authentication, the client is returned to the HTTP connection.**

If there are clients that are not HTTP-enabled connected to the captive portal interface, you can allow them to bypass captive portal authentication by adding their MAC address to an authentication whitelist. (If the MAC address has already been learned on the interface, you must clear it using the **clear captive-portal interface** ***interface-name***) before adding it to the whitelist.)

When the user is authenticated by the RADIUS server, any per-user policies (attributes) associated with that user are also sent to the switch.

Limitations of Captive Portal

Captive portal on MX Series routers has the following limitations:

- The captive portal interface must be configured for **family ethernet-switching** and set to port mode access. The VLAN must be configured with a routed VLAN interface (RVI).
- The DHCP gateway IP address for the switch must be configured as the IP address of the routed VLAN interface.
- Captive portal does not support dynamic assignment of VLANs downloaded from the RADIUS server.
- If the user is idle for more than about 5 minutes and there is no traffic passed, the user is required to log back in to the captive portal.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, captive portal authentication (hereafter referred to as captive portal) allows you to authenticate users on MX Series routers by redirecting Web browser requests to a login page that requires users to input a username and password before they are allowed access to the network.

Understanding Authentication Session Timeout on MX Series Routers

Starting with Junos OS Release 14.2, you can specify authentication session timeout values for captive portal authentication sessions and 802.1X and MAC RADIUS authentication sessions.

For captive portal authentication, the length of the session depends on the value configured for the **session-expiry** statement. The remainder of this topic pertains only to 802.1X and MAC RADIUS authentication sessions.

For 802.1X and MAC RADIUS authentication sessions, the timeout of the session depends on the value of **reauthentication interval** for **dot1x authentication**. The authentication session might also end when the MAC table aging time expires because, unless you configure it not to, the session is removed from the authentication session table when the MAC address is removed from the Ethernet switching table.

Information about each 802.1X and MAC RADIUS authentication session—including the associated interfaces and VLANs for each MAC address that is authenticated by 802.1X authentication or MAC RADIUS authentication—is stored in the authentication session table. The authentication session table is tied to the Ethernet switching table (also called the MAC table). Each time the switch detects traffic from a MAC address, it updates the timestamp for that network node in the Ethernet switching table. A timer on the switch periodically checks the timestamp and if its value exceeds the user-configured **mac-table-aging-time** value, the switch removes the MAC address from the Ethernet switching table. When a MAC address ages out of the Ethernet switching table, the entry for that MAC address is also removed from the authentication database, with the result that the session ends.

You can control variables affecting timeout of authentication sessions in the following ways:

- Set the authentication session timeout on all interfaces or on selected interfaces using the **reauthentication** statement.
- Disassociate the authentication session table from the Ethernet switching table using the **no-mac-table-binding** statement. This setting prevents the termination of the authentication session when the associated MAC address ages out of the Ethernet switching table.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can specify authentication session timeout values for captive portal authentication sessions and 802.1X and MAC RADIUS authentication sessions.

Authentication Process Flow for MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, you can control access to your network through an MX Series router by using several different authentication methods—including 802.1X, MAC RADIUS, or captive portal.

Figure 3 on page 64 illustrates the authentication process:

Figure 3: Authentication Process Flow for an MX Series Router

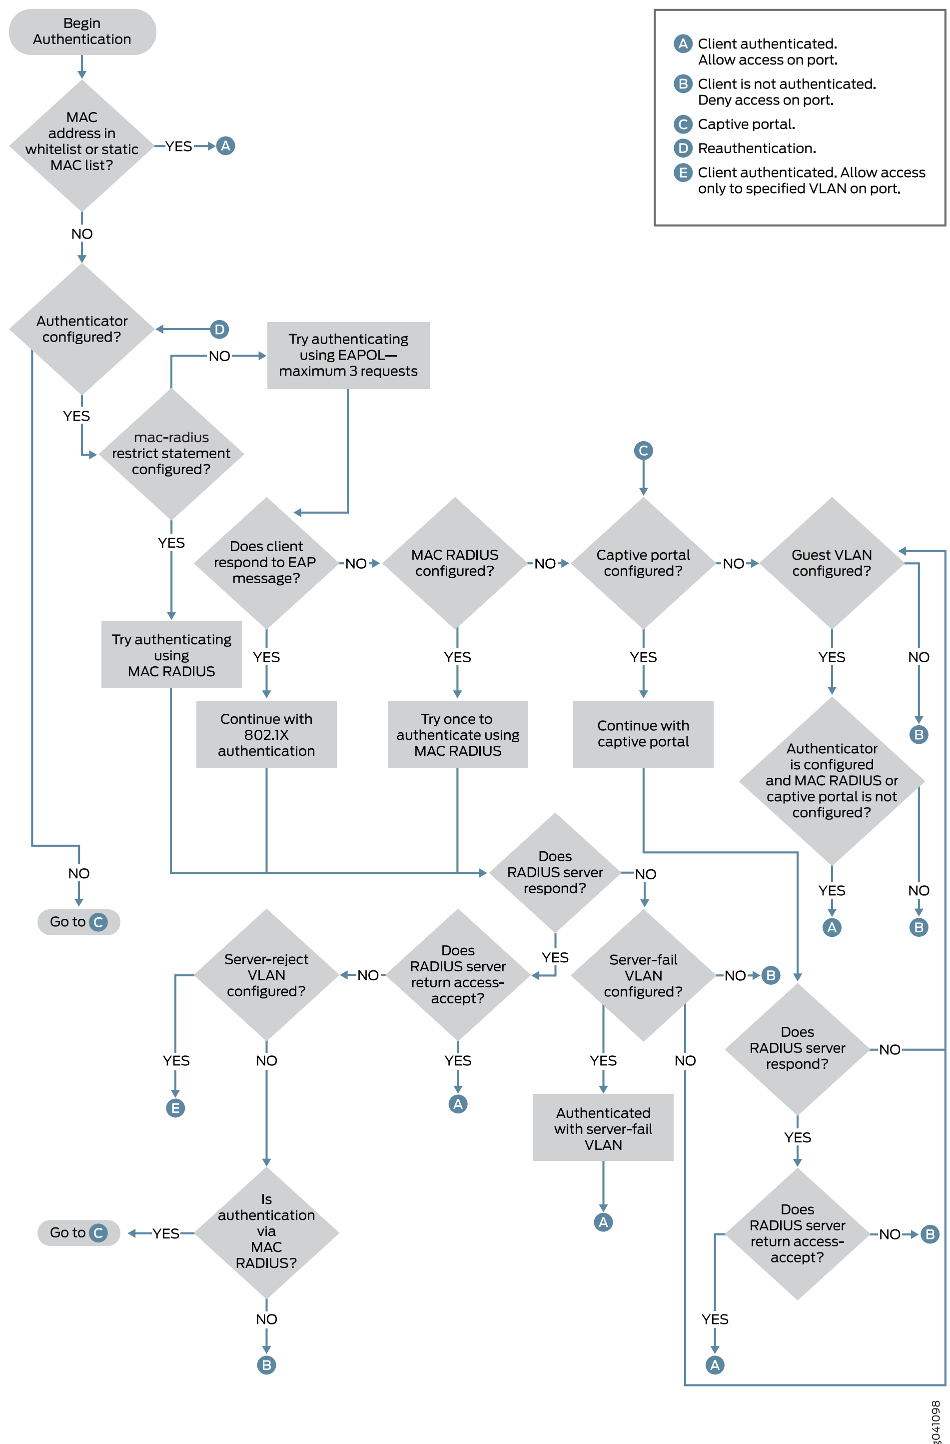

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can control access to your network through an MX Series router by using several different authentication methods—including 802.1X, MAC RADIUS, or captive portal.

Specifying RADIUS Server Connections on an MX Series Router in Enhanced LAN Mode

IEEE 802.1X and MAC RADIUS authentication both provide network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from devices at the interface until the supplicant's credentials or MAC address are presented and matched on the *authentication server* (a RADIUS server). When the supplicant is authenticated, the router stops blocking access and opens the interface to the supplicant.

Starting with Junos OS Release 14.2, to use 802.1X or MAC RADIUS authentication, you must specify the connections on the router for each RADIUS server to which you will connect.

To configure a RADIUS server on the router:

1. Define the IP address of the RADIUS server, the RADIUS server authentication port number, and the secret password. You can define more than one RADIUS server. The secret password on the router must match the secret password on the server:

   ```
   [edit access]
   user@router# set radius-server 10.0.0.100 port 1812 secret abc
   ```

   NOTE: **Specifying the authentication port is optional, and port 1812 is the default. However, we recommend that you configure it in order to avoid confusion as some RADIUS servers might refer to an older default.**

2. (Optional) Specify the IP address by which the router is identified by the RADIUS server. If you do not specify this, the RADIUS server uses the address of the interface sending the RADIUS request. We recommend that you specify this IP address because if the request gets diverted on an alternate route to the RADIUS server, the interface relaying the request might not be an interface on the router.

   ```
   [edit access]
   user@router# set radius-server source-address 10.93.14.100
   ```

3. Configure the authentication order, making **radius** the first method of authentication:

   ```
   [edit access]
   user@router# set profile profile1 authentication-order radius
   ```

4. Create a profile and specify the list of RADIUS servers to be associated with the profile. For example, you might choose to group your RADIUS servers geographically by city.

This feature enables easy modification whenever you want to change to a different sent of authentication servers.

```
[edit access profile]
user@router# set atlanta radius authentication-server 10.0.0.100 10.2.14.200
```

5. Specify the group of servers to be used for 802.1X or MAC RADIUS authentication by identifying the profile name:

```
[edit access profile]
user@router# set protocols authentication-access-control authentication-profile-name
denver
```

6. Configure the IP address of the MX Series router in the list of clients on the RADIUS server. For specifics on configuring the RADIUS server, consult the documentation for your server.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, to use 802.1X or MAC RADIUS authentication, you must specify the connections on the router for each RADIUS server to which you will connect.

Configuring Captive Portal Authentication on MX Series Routers in Enhanced LAN Mode

NOTE: This example uses Junos OS for MX240, MX480, and MX960 routers with support for the Enhanced LAN mode configuration style. If your router does not run MX-LAN mode, you cannot configure port-based authentication settings in the same manner as described in this section. If you remove the network-services lan statement at the [edit chassis] hierarchy level, the system does not run in MX-LAN mode. Therefore, all of the settings that are supported outside of the MX-LAN mode are displayed and are available for definition in the CLI interface. In such a scenario, you must use the statements at the **[edit protocols dot1x]** hierarchy level to configure 802.1x and MAC RADIUS authentication, and the options at the **[edit services captive-portal]** hierarchy level to configure captive portal authentication. In MX-LAN mode, you can configure all the port-based network access control methodologies using the statements at the **[edit protocols authentication-access-control]** hierarchy level.

Starting with Junos OS Release 14.2, configure captive portal authentication (hereafter referred to as captive portal) on an MX Series router so that users connected to the router are authenticated before being allowed to access the network. When the user requests a webpage, a login page is displayed that requires the user to input a username and password. Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

Before you begin, be sure you have:

- Performed basic bridging and VLAN configuration on the router.
- Generated an SSL certificate and installed it on the router.
- Configured basic access between the MX Series router and the RADIUS server.
- Designed your captive portal login page.

This topic includes the following tasks:

- Configuring Secure Access for Captive Portal on page 67
- Enabling an Interface for Captive Portal on page 67
- Configuring Bypass of Captive Portal Authentication on page 67

Configuring Secure Access for Captive Portal

To configure secure access for captive portal:

1. Associate the security certificate with the Web server and enable HTTPS on the router:

   ```
   [edit]
   user@router# set system services web-management https local-certificate my-signed-cert
   ```

   NOTE: You can enable HTTP instead of HTTPS, but we recommend HTTPS for security purposes.

2. Configure captive portal to use HTTPS:

   ```
   [edit]
   user@router# set protocols custom-options-captive-portal secure-authentication https
   ```

Enabling an Interface for Captive Portal

To enable an interface for use with captive portal authentication:

```
[edit]
user@router# set authentication-access-control interface ge-0/0/10
```

Configuring Bypass of Captive Portal Authentication

You can allow specific clients to bypass captive portal authentication:

```
[edit]
user@router# set authentication-access-control static 00:10:12:e0:28:22
```

NOTE: Optionally, you can use **set authentication-access-control static 00:10:12:e0:28:22 interface ge-0/0/10.0** to limit the scope to the interface.

NOTE: If the client is already attached to the router, you must clear its MAC address from the captive portal authentication by using the **clear captive-portal mac-address** ***session-mac-addr*** command after adding its MAC address to the whitelist. Otherwise the new entry for the MAC address will not be added to the Ethernet switching table and the authentication bypass will not be allowed.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, configure captive portal authentication (hereafter referred to as captive portal) on an MX Series router so that users connected to the router are authenticated before being allowed to access the network.

Designing a Captive Portal Authentication Login Page on an MX Series Router

Starting with Junos OS Release 14.2, you can set up captive portal authentication on your switch to redirect all Web browser requests to a login page that requires the user to input a username and password before they are allowed access.Upon successful authentication, the user is allowed access to the network and redirected to the original page requested.

Junos OS provides a customizable template for the captive portal window that allows you to easily design and modify the look of the captive portal login page. You can modify the design elements of the template to change the look of your captive portal login page and to add instructions or information to the page. You can also modify any of the design elements of a captive portal login page.

The first screen displayed before the captive login page requires the user to read the “Terms and Conditions of Use”. By clicking the Agree button, the user can access the captive portal login page.

Figure 4 on page 69 shows an example of a captive portal login page:

Figure 4: Example of a Captive Portal Login Page

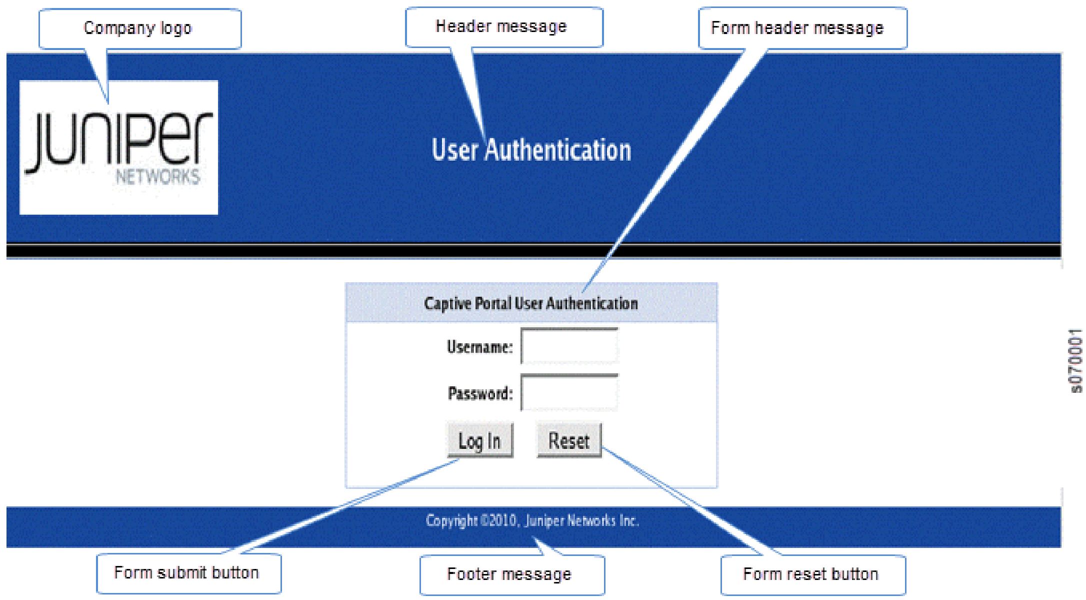

Table 3 on page 69 summarizes the configurable elements of a captive portal login page.

Table 3: Configurable Elements of a Captive Portal Login Page

Element	CLI Statement	Description
Footer background color	**footer-bgcolor** ***hex-color***	The HTML hexadecimal code for the background color of the captive portal login page footer.
Footer message	**footer-message** ***text-string***	Text displayed in the footer of the captive portal login page. You can include copyright information, links, and additional information such as help instructions, legal notices, or a privacy policy The default text shown in the footer is **Copyright @2010, Juniper Networks Inc.**
Footer text color	**footer- text-color** ***color***	Color of the text in the footer. The default color is white.
Form header background color	**form-header-bgcolor** ***hex-color***	The HTML hexadecimal code for the background color of the header bar across the top of the form area of the captive portal login page.
Form header message	**form-header-message** ***text-string***	Text displayed in the header of the captive portal login page. The default text is **Captive Portal User Authentication**
Form header text color	**form-header- text- color** *color*	Color of the text in the form header. The default color is black.
Form reset button label	**form-reset-label** ***label-name***	Using the **Reset** button, the user can clear the username and password fields on the form.
Form submit button label	**form-submit-label** ***label-name***	Using the **Login** button, the user can submit the login information.

Table 3: Configurable Elements of a Captive Portal Login Page (continued)

Element	CLI Statement	Description
Header background color	**header-bgcolor** *hex-color*	The HTML hexadecimal code for the background color of the captive portal login page header.
Header logo	**header-logo** *filename*	Filename of the file containing the image of the logo that you want to appear in the header of the captive portal login page. The image file can be in GIF, JPEG, or PNG format You can upload a logo image file to the switch. Copy the logo to the /var/tmp directory on the switch (during commit, the files are saved to persistent locations). If you do not specify a logo image, the Juniper Networks logo is displayed.
Header message	**header-message** *text-string*	Text displayed in the page header. The default text is **User Authentication**.
Header text color	**header-text- color***color*	Color of the text in the header. The default color is white.
Post-authentication URL	**post-authentication-url** *url*	URL to which the users are directed on successful authentication. By default, users are directed to the page they had originally requested.

To design the captive portal login page:

1. (Optional) Upload your logo image file to the switch:

   `user@router>` **file copy ftp://username:prompt@ftp.hostname.net/var/tmp/my-logo.jpeg**

2. Configure the custom options to specify the background colors and text displayed in the captive portal page:

   `[edit protocols]`
   `user@router#` **set captive-portal-custom-options header-bgcolor #006600**
   set captive-portal-custom-options header-message “Welcome to Our Network”
   set captive-portal-custom-options banner-message “Please enter your username and password”.The banner displays the message ”XXXXXXX” by default. The user can modify this message.
   set custom-options footer-message “Copyright ©2010, Our Network”

Now you can commit the configuration.

NOTE: For the custom options that you do not specify, the default value is used.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can set up captive portal authentication on your switch to redirect all Web browser requests to a login page that requires the user to input a username and password before they are allowed access.

Configuring Static MAC Bypass of Authentication on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, you can configure a static MAC bypass list (sometimes called the exclusion list) on the switch to specify MAC addresses of devices allowed access to the LAN without 802.1X or MAC RADIUS authentication requests to the RADIUS server.

To configure the static MAC bypass list:

- Specify a MAC address to bypass authentication:

  ```
  [edit protocols authentication-access-control]
  user@router# set static 00:04:0f:fd:ac:fe
  ```

- Configure a supplicant to bypass authentication if connected through a particular interface:

  ```
  [edit protocols authentication-access-control]
  user@router# set static 00:04:0f:fd:ac:fe interface ge-0/0/5
  ```

- You can configure a supplicant to be moved to a specific VLAN after it is authenticated:

  ```
  [edit protocols authentication-access-control]
  user@router# set static 00:04:0f:fd:ac:fe interface ge-0/0/5 vlan-assignment default-vlan
  ```

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can configure a static MAC bypass list (sometimes called the exclusion list) on the switch to specify MAC addresses of devices allowed access to the LAN without 802.1X or MAC RADIUS authentication requests to the RADIUS server.

Controlling Authentication Session Timeouts on an MX Series Router in Enhanced LAN Mode

Starting with Junos OS Release 14.2, for 802.1X and MAC RADIUS authentication sessions, you can specify authentication session timeout values using the **reauthentication** statement.

The session might also end when the MAC table aging time expires, because the session is removed from the authentication session table when the MAC address is removed from the Ethernet switching table. In order to prevent the session from being removed from the authentication session table, you must disassociate the authentication table from the Ethernet switching table using the **no-mac-table-binding** statement.

Before you begin:

- Specify the RADIUS server or servers to be used as the authentication server.
- Configure 802.1X authentication on the router.

To configure the authentication session time on all interfaces:

```
[edit]
user@router# set protocols authentication-access-control interface all dot1x reauthentication seconds ;
```

To configure the authentication session time on a single interface:

```
[edit]
user@router# set protocols authentication-access-control interface interface-name dot1x reauthentication seconds ;
```

To disable removal of authentication sessions from the authentication session table when a MAC address ages out of the Ethernet switching table, remove the binding of the authentication table to the Ethernet switching table.

To remove the binding on all interfaces:

```
[edit]
user@router# set protocols authentication-access-control no-mac-table-binding interface all;
```

To remove the binding on a single interface:

```
[edit]
user@router# set protocols authentication-access-control no-mac-table-binding interface interface-name;
```

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, for 802.1X and MAC RADIUS authentication sessions, you can specify authentication session timeout values using the **reauthentication** statement.

Related Documentation

Configuring MAC RADIUS Authentication on MX Series Routers in Enhanced LAN Mode

Starting with Junos OS Release 14.2, you can permit devices that are not 802.1X-enabled LAN access by configuring MAC RADIUS authentication on the MX Series router interfaces to which the hosts are connected.

NOTE: **You can also allow non-802.1X-enabled devices to access the LAN by configuring their MAC address for static MAC bypass of authentication.**

You can configure MAC RADIUS authentication on an interface that also allows 802.1X authentication, or you can configure either authentication method alone.

If both MAC RADIUS and 802.1X authentication are enabled on the interface, the router first sends the host three EAPOL requests to the host. If there is no response from the host, the router sends the host's MAC address to the RADIUS server to check whether it is a permitted MAC address. If the MAC address is configured as permitted on the RADIUS server, the RADIUS server sends a message to the router that the MAC address is a permitted address, and the router opens LAN access to the nonresponsive host on the interface to which it is connected.

If MAC RADIUS authentication is configured on the interface but 802.1X authentication is not (by using the **mac-radius restrict** option), the router attempts to authenticate the MAC address with the RADIUS server without delaying by attempting 802.1X authentication first.

Before you configure MAC RADIUS authentication, be sure you have:

- Configured basic access between the MX Series router and the RADIUS server.
- Configured MX240, MX480, and MX960 routers to function in enhanced LAN mode by entering the **network-services lan** statement at the **[edit chassis]** hierarchy level.

To configure MAC RADIUS authentication using the CLI:

- On the router, configure the interfaces to which the nonresponsive hosts are attached for MAC RADIUS authentication, and add the **restrict** qualifier for interface **ge-0/0/20** to have it use only MAC RADIUS authentication:

```
[edit]
user@router# set protocols authentication-access-control interface ge-0/0/19 dot1x mac-radius
user@router# set protocols authentication-access-control interface ge-0/0/20 dot1x mac-radius restrict
```

- On a RADIUS authentication server, create user profiles for each nonresponsive host using the MAC address (without colons) of the nonresponsive host as the username and password (here, the MAC addresses are **00:04:0f:fd:ac:fe** and **00:04:ae:cd:23:5f**):

```
[root@freeradius]#
edit /etc/raddb
vi users
00040ffdacfe Auth-type:=Local, User-Password = "00040ffdacfe"
0004aecd235f Auth-type:=Local, User-Password = "0004aecd235f"
```

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can permit devices that are not 802.1X-enabled LAN access by configuring MAC RADIUS authentication on the MX Series router interfaces to which the hosts are connected.

Related Documentation

Example: Configuring MAC RADIUS Authentication on an MX Series Router

Starting with Junos OS Release 14.2 to permit hosts that are not 802.1X-enabled to access the LAN, you can configure MAC RADIUS authentication on the router interfaces to which the non-802.1X-enabled hosts are connected. When MAC RADIUS authentication is configured, the router will attempt to authenticate the host with the RADIUS server using the host's MAC address.

This example describes how to configure MAC RADIUS authentication for two non-802.1X-enabled hosts:

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.2 or later for MX240, MX480, or MX960 routers running in enhanced LAN mode.
- An MX Series router acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- A RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the router, be sure you have:

- Configured enhanced LAN mode on the router.
- Performed basic bridging and VLAN configuration on the router.
- Configured users on the RADIUS authentication server.

Overview and Topology

IEEE 802.1X Port-Based Network Access Control (PNAC) authenticates and permits devices access to a LAN if the devices can communicate with the router using the 802.1X protocol (are 802.1X-enabled). To permit non-802.1X-enabled end devices to access the LAN, you can configure MAC RADIUS authentication on the interfaces to which the end devices are connected. When the MAC address of the end device appears on the interface, the router consults the RADIUS server to check whether it is a permitted MAC address. If the MAC address of the end device is configured as permitted on the RADIUS server, the router opens LAN access to the end device.

You can configure both MAC RADIUS authentication and 802.1X authentication methods on an interface configured for multiple supplicants. Additionally, if an interface is only connected to a non-802.1X-enabled host, you can enable MAC RADIUS and not enable 802.1X authentication using the **mac-radius restrict** option, and thus avoid the delay that occurs while the router determines that the device is does not respond to EAP messages.

Two printers are connected to an MX Series router over interfaces, ge-0/0/19 and ge-0/0/20.

Table 4 on page 76 shows the components in the example for MAC RADIUS authentication.

Table 4: Components of the MAC RADIUS Authentication Configuration Topology

Property	Settings
Router hardware	Ports (ge-0/0/0 through ge-0/0/23)
VLAN name	sales

Table 4: Components of the MAC RADIUS Authentication Configuration Topology (continued)

Property	Settings
Connections to printers	ge-0/0/19, MAC address 00040ffdacfe ge-0/0/20, MAC address 0004aecd235f
RADIUS server	Connected to the router on interface **ge-0/0/10**

The printer with the MAC address 00040ffdacfe is connected to access interface ge-0/0/19. A second printer with the MAC address 0004aecd235f is connected to access interface ge-0/0/20. In this example, both interfaces are configured for MAC RADIUS authentication on the router, and the MAC addresses (without colons) of both printers are configured on the RADIUS server. Interface ge-0/0/20 is configured to eliminate the normal delay while the router attempts 802.1X authentication; MAC RADIUS authentication is enabled and 802.1X authentication is disabled using the **mac radius restrict** option.

Configuration

CLI Quick Configuration

To quickly configure MAC RADIUS authentication, copy the following commands and paste them into the router terminal window:

```
[edit]
set protocols authentication-access-control interface ge-0/0/19 dot1x mac-radius
set protocols authentication-access-control authenticator interface ge-0/0/20 dot1x mac-radius
restrict
```

NOTE: You must also configure the two MAC addresses as usernames and passwords on the RADIUS server, as is done in step 2 of the Step-by-Step Procedure.

Step-by-Step Procedure

Configure MAC RADIUS authentication on the router and on the RADIUS server:

1. On the router, configure the interfaces to which the printers are attached for MAC RADIUS authentication, and configure the **restrict** option on interface **ge-0/0/20**, so that only MAC RADIUS authentication is used:

```
[edit]
user@router# set protocols authentication-access-control interface ge-0/0/19 dot1x
mac-radius
user@router# set protocols authentication-access-control authenticator interface
ge-0/0/20 dot1x mac-radius restrict
```

2. On the RADIUS server, configure the MAC addresses **00040ffdacfe** and **0004aecd235f** as usernames and passwords:

```
[root@freeradius]#
edit /etc/raddb
```

```
vi users
00040ffdacfe Auth-type:=EAP, User-Password = "00040ffdacfe"
0004aecd235f Auth-type:=EAP, User-Password = "0004aecd235f"
```

Results Display the results of the configuration on the router:

```
user@router> show configuration
protocols {
  authentication-access-control {
    authentication-profile-name profile52;
    interface {
      ge-0/0/19.0 {
        dot1x {
          mac-radius;
        }
      }
      ge-0/0/20.0 {
        dot1x {
          mac-radius {
            restrict;
          }
        }
      }
    }
  }
}
```

Verification

Verify that the supplicants are authenticated:

- Verifying That the Supplicants Are Authenticated on page 78

Verifying That the Supplicants Are Authenticated

Purpose After supplicants are configured for MAC RADIUS authentication on the router and on the RADIUS server, verify that they are authenticated and display the method of authentication:

Action Display information about 802.1X-configured interfaces **ge-0/0/19** and **ge-0/0/20**:

```
user@router> show dot1x interface ge-0/0/19.0 detail
ge-0/0/19.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user101, 00:04:0f:fd:ac:fe
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

```
user@router> show dot1x interface ge-0/0/20.0 detail
ge-0/0/20.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Enabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user102, 00:04:ae:cd:23:5f
      Operational state: Authenticated
      Authentcation method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

Meaning The sample output from the **show dot1x interface detail** command displays the MAC address of the connected end device in the **Supplicant** field. On interface ge-0/0/19, the MAC address is **00:04:0f:fd:ac:fe**, which is the MAC address of the first printer configured for MAC RADIUS authentication. The **Authentication method** field displays the authentication method as **MAC Radius**. On interface **ge-0/0/20**, the MAC address is

00:04:ae:cd:23:5f, which is the MAC address of the second printer configured for MAC RADIUS authentication. The **Authentication method** field displays the authentication method as **MAC Radius**.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2 to permit hosts that are not 802.1X-enabled to access the LAN, you can configure MAC RADIUS authentication on the router interfaces to which the non-802.1X-enabled hosts are connected.

Example: Setting Up Captive Portal Authentication on an MX Series Router

Starting with Junos OS Release 14.2, you can set up captive portal authentication (hereafter referred to as captive portal) on a router to redirect Web browser requests to a login page that requires the user to input a username and password.Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

This example describes how to set up captive portal on an MX Series router:

Requirements

This example uses the following hardware and software components:

- An MX Series router that supports captive portal
- Junos OS Release 14.2 or later for MX Series routers

Before you begin, be sure you have:

- Performed basic bridging and VLAN configuration on the router.
- Generated an SSL certificate and installed it on the router.
- Configured basic access between the MX Series router and the RADIUS server.
- Designed your captive portal login page. .

Overview and Topology

This example shows the configuration required on the router to enable captive portal on an interface. To permit a printer connected to the captive portal interface to access the LAN without going through captive portal, add its MAC address to the authentication whitelist. The MAC addresses in this list are permitted access on the interface without captive portal.

The topology for this example consists of one MX Series router connected to a RADIUS authentication server. One interface on the router is configured for captive portal. In this example, the interface is configured in multiple supplicant mode.

Configuration

To configure captive portal on your router:

CLI Quick Configuration

To quickly configure captive portal on the router after completing the tasks in the Requirements section, copy the following commands and paste them into the router terminal window:

```
[edit]
set system services web-management http
set system services web-management https local-certificate my-signed-cert
set protocols captive-portal-custom-options secure-authentication https
set protocols authentication-access-control interface ge-0/0/10.0 supplicant multiple
set protocols authentication-access-control static 00:10:12:e0:28:22
set protocols captive-portal-custom-options post-authentication-url
http://www.my-home-page.com
```

Step-by-Step Procedure

To configure captive portal on the router:

1. Enable HTTP access on the router:

   ```
   [edit]
   user@router# set system services web-management http
   ```

2. To create a secure channel for Web access to the router, configure captive portal for HTTPS:

   NOTE: You can enable HTTP without enabling HTTPS, but we recommend HTTPS for security purposes.

   a. Associate the security certificate with the Web server and enable HTTPS access on the router:

   ```
   [edit]
   user@router# set system services web-management https local-certificate
   my-signed-cert
   ```

   b. Configure captive portal to use HTTPS:

   ```
   [edit]
   user@router# set protocols captive-portal-custom-options secure-authentication
   https
   ```

3. Enable an interface for captive portal:

```
[edit]
user@router# set protocols authentication-access-control interface ge-0/0/10.0 supplicant multiple
```

4. (Optional) Allow specific clients to bypass captive portal:

   NOTE: If the client is already attached to the router, you must clear its MAC address from the captive portal authentication by using the **clear captive-portal mac-address** *mac-address* command after adding its MAC address to the whitelist. Otherwise the new entry for the MAC address will not be added to the Ethernet routering table and authentication bypass will not be allowed.

```
[edit]
user@router# set protocols authentication-access-control static 00:10:12:e0:28:22
```

   NOTE: Optionally, you can use **set ethernet-switching-options authentication-whitelist 00:10:12:e0:28:22 interface ge-0/0/10.0** to limit the scope to the interface.

5. (Optional) To redirect clients to a specified page rather than the page they originally requested, configure the post-authentication URL:

```
[edit services captive-portal]
user@router# set protocols captive-portal-custom-options post-authentication-url http://www.my-home-page.com
```

Results Display the results of the configuration:

```
[edit]
user@router> show
system {
  services {
    web-management {
      http;
      https {
        local-certificate my-signed-cert;
      }
    }
  }
}
security {
  certificates {
    local {
      my-signed-cert {
```



```
                "-----BEGIN RSA PRIVATE KEY-----\nMIICXwIBAAKBgQDk8sUggnXdDUmr7T
                  vLv63yJq/LRpDASfIDZlX3z9ZDe1Kfk5C9\nr/tkyvzv
                  ...
                  Pt5YmvWDoGo0mSjoE/liH0BqYdh9YGqv3T2IEUfflSTQQHEOShS0ogWDHF\
                  nnyOb1O/vQtjk20X9NVQg JHBwidssY9eRp\n-----END CERTIFICATE-----\n";
                  ## SECRET-DATA
            }
        }
    }
}
protocols {
    authentication-access-control {
    static 00:10:12:e0:28:22/48;
        interface {
            ge-0/0/10.0 {
                supplicant multiple;
            }
        }
    custom-captive-portal-options {
        secure-authentication https;
        post-authentication-url http://www.my-home-page.com;
    }
}
```

Verification

To confirm that captive portal is configured and working properly, perform these tasks:

- Verifying That Captive Portal Is Enabled on the Interface on page 83
- Verify That Captive Portal Is Working Correctly on page 83

Verifying That Captive Portal Is Enabled on the Interface

Purpose Verify that captive portal is configured on interface ge-0/0/10.

Action Use the operational mode command **show captive-portal interface *interface-name* detail**:

```
user@router> show captive-portal interface ge-0/0/10.0 detail
ge-0/0/10.0
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
```

Meaning The output confirms that captive portal is configured on interface ge-0/0/10 with the default settings for number of retries, quiet period, CP session timeout, and server timeout.

Verify That Captive Portal Is Working Correctly

Purpose Verify that captive portal is working on the router.

Action Connect a client to interface ge-0/0/10. From the client, open a Web browser and request a webpage. The captive portal login page that you designed should be displayed. After you enter your login information and are authenticated against the RADIUS server, the Web browser should display either the page you requested or the post-authentication URL that you configured.

Troubleshooting

To troubleshoot captive portal, perform these tasks:

Troubleshooting Captive Portal

Problem The router does not return the captive portal login page when a user connected to a captive portal interface on the router requests a Web page.

Solution You can examine the ARP, DHCP, HTTPS, and DNS counters—if one or more of these counters are not incrementing, this provides an indication of where the problem lies. For example, if the client cannot get an IP address, check the router interface to determine whether the DHCP counter is incrementing—if the counter increments, the DHCP packet was received by the router.

```
user@router> show captive-portal firewall ge-0/0/10.0
ge-0/0/10.0
  Filter name: dot1x_ge-0/0/10
Counters:
Name                          Bytes              Packets
dot1x_ge-0/0/10_CP_arp         7616                  119
dot1x_ge-0/0/10_CP_dhcp           0                    0
dot1x_ge-0/0/10_CP_http           0                    0
dot1x_ge-0/0/10_CP_https          0                    0
dot1x_ge-0/0/10_CP_t_dns          0                    0
dot1x_ge-0/0/10_CP_u_dns          0                    0
```

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can set up captive portal authentication (hereafter referred to as captive portal) on a router to redirect Web browser requests to a login page that requires the user to input a username and password.

Example: Connecting a RADIUS Server for 802.1X to an MX Series Router

802.1X is the IEEE standard for Port-Based Network Access Control (PNAC). You use 802.1X to control network access. Only users and devices providing credentials that have been verified against a user database are allowed access to the network. Starting with Junos OS Release 14.2, you can use a RADIUS server as the user database for 802.1X authentication, as well as for MAC RADIUS authentication.

This example describes how to connect a RADIUS server to an MX Series router, and configure it for 802.1X:

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.2 or later for MX240, MX480, or MX960 routers running in enhanced LAN mode and Junos OS Release 14.2R3 for all other routers.
- One router acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the router, be sure you have:

- Configured enhanced LAN mode on the router.
- Performed basic bridging and VLAN configuration on the router.
- Configured users on the RADIUS authentication server.

Overview and Topology

The MX Series router acts as an authenticator Port Access Entity (PAE). It blocks all traffic and acts as a control gate until the supplicant (client) is authenticated by the server. All other users and devices are denied access.

Consider an MX Series router that functions as an authenticator port. It is connected using the interface, ge-0/0/10, over the IP network to a RADIUS server. The router is also linked to a conference room using the interface, ge-0/0/1, to a printer using the interface, ge-0/0/20, to a hub using the interface, ge-0/0/8, and to two supplicants or clients over interfaces, ge-0/0/2 and ge-0/0/9 respectively.

Table 5: Components of the Topology

Property	Settings
Router hardware	MX Series router
VLAN name	**default**
One RADIUS server	Backend database with an address of **10.0.0.100** connected to the switch at port **ge-0/0/10**

In this example, connect the RADIUS server to access port **ge-0/0/10** on the MX Series router. The switch acts as the authenticator and forwards credentials from the supplicant to the user database on the RADIUS server. You must configure connectivity between the MX Series router and the RADIUS server by specifying the address of the server and configuring the secret password. This information is configured in an access profile on the switch.

Configuration

CLI Quick Configuration

To quickly connect the RADIUS server to the switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set access radius-server 10.0.0.100 secret juniper
set access radius-server 10.0.0.200 secret juniper
set access profile profile1 authentication-order radius
set access profile profile1 radius authentication-server [10.0.0.100 10.0.0.200]
```

Step-by-Step Procedure

To connect the RADIUS server to the switch:

1. Define the address of the servers, and configure the secret password. The secret password on the switch must match the secret password on the server:

   ```
   [edit]
   user@switch# set access radius-server 10.0.0.100 secret juniper
   user@switch# set access radius-server 10.0.0.200 secret juniper
   ```

2. Configure the authentication order, making **radius** the first method of authentication:

   ```
   [edit]
   user@switch# set access profile profile1 authentication-order radius
   ```

3. Configure a list of server IP addresses to be tried in order to authenticate the supplicant:

   ```
   [edit]
   user@switch# set access profile profile1 radius authentication-server [10.0.0.100
   10.0.0.200]
   ```

Results

Display the results of the configuration:

```
user@switch> show configuration access
radius-server {
  10.0.0.100
    port 1812;
    secret "$9$qPT3ApBSrv69rvWLVb.P5"; ## SECRET-DATA
    }
  }
  profile profile1{
    authentication-order radius;
    radius {
      authentication-server 10.0.0.100 10.0.0.200;
```

```
        }
      }
    }
```

Verification

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

Verify That the Switch and RADIUS Server are Properly Connected

Purpose Verify that the RADIUS server is connected to the switch on the specified port.

Action Ping the RADIUS server to verify the connection between the switch and the server:

```
user@switch> ping 10.0.0.100
PING 10.0.0.100 (10.0.0.100): 56 data bytes
64 bytes from 10.93.15.218: icmp_seq=0 ttl=64 time=9.734 ms
64 bytes from 10.93.15.218: icmp_seq=1 ttl=64 time=0.228 ms
```

Meaning ICMP echo request packets are sent from the switch to the target server at 10.0.0.100 to test whether it is reachable across the IP network. ICMP echo responses are being returned from the server, verifying that the switch and the server are connected.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can use a RADIUS server as the user database for 802.1X authentication, as well as for MAC RADIUS authentication.

Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on an MX Series Router

Starting with Junos OS Release 14.2, 802.1X on MX Series routers provides LAN access to users who do not have credentials in the RADIUS database.These users, referred to as guests, are authenticated and typically provided with access to the Internet.

This example describes how to create a guest VLAN and configure 802.1X authentication for it.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.2 or later for MX240, MX480, or MX960 routers running in enhanced LAN mode.
- One router acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the router, be sure you have:

- Configured enhanced LAN mode on the router.
- Performed basic bridging and VLAN configuration on the router.
- Configured users on the RADIUS authentication server.

Overview and Topology

The MX Series router acts as an authenticator Port Access Entity (PAE). It blocks all traffic and acts as a control gate until the supplicant (client) is authenticated by the server. All other users and devices are denied access.

Consider an MX Series router that functions as an authenticator port. It is connected using the interface, ge-0/0/10, over the IP network to a RADIUS server. The router is also linked to a conference room using the interface, ge-0/0/1, to a printer using the interface, ge-0/0/20, to a hub using the interface, ge-0/0/8, and to two supplicants or clients over interfaces, ge-0/0/2 and ge-0/0/9 respectively.

Table 6: Components of the Topology

Property	Settings
Router hardware	MX Series router
VLAN name	**default**
One RADIUS server	Backend database with an address of **10.0.0.100** connected to the switch at port **ge-0/0/10**

In this example, access interface **ge-0/0/1** provides LAN connectivity in the conference room. Configure this access interface to provide LAN connectivity to visitors in the conference room who are not authenticated by the corporate VLAN.

Configuration of a Guest VLAN That Includes 802.1X Authentication

CLI Quick Configuration

To quickly configure a guest VLAN, with 802.1X authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans guest-vlan vlan-id 300
set protocols authentication-access-control interface all dot1x guest-vlan guest-vlan
```

Step-by-Step Procedure

To configure a guest VLAN that includes 802.1X authentication on an EX Series switch:

1. Configure the VLAN ID for the guest VLAN:

   ```
   [edit]
   user@switch# set bridge-domains guest-vlan vlan-id 300
   ```

2. Configure the guest VLAN under **dot1x** protocols:

   ```
   [edit]
   user@switch# set protocols authentication-access-control interface all dot1x guest-vlan guest-vlan
   ```

Results

Check the results of the configuration:

```
user@switch> show configuration
protocols {
  authentication-access-control {
    interface {
      all {
        dot1x {
          guest-vlan {
            guest-vlan;
          }
        }
      }
    }
    }
  }
}
bridge-domains {
  guest-vlan {
    vlan-id 300;
  }
}
```

Verification

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

- Verifying That the Guest VLAN is Configured on page 90

Verifying That the Guest VLAN is Configured

Purpose Verify that the guest VLAN is created and that an interface has failed authentication and been moved to the guest VLAN.

Action Use the operational mode commands:

user@switch> **show bridge-domain**

```
Instance                     Bridging Domain          Type
            Primary Table                                             Active
vs1                          dynamic                   bridge
            bridge.0                                                  2
vs1                          guest                     bridge
            bridge.0                                                  0
vs1                          guest-vlan                bridge
            bridge.0                                                  0
vs1                          vlan_dyn                  bridge
            bridge.0                                                  0
```

user@switch> **show dot1x interface ge-0/0/1.0 detail**

```
ge-0/0/1.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: guest-vlan
  Number of connected supplicants: 1
    Supplicant: user1, 00:00:00:00:13:23
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

Meaning The output from the **show bridge domain** command shows **guest-vlan** as the name of the VLAN and the VLAN ID as **300**.

The output from the **show dot1x interface ge-0/0/1.0 detail** command displays the **Guest VLAN membership** field, indicating that a supplicant at this interface failed 802.1X authentication and was passed through to the **guest-vlan**.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, 802.1X on MX Series routers provides LAN access to users who do not have credentials in the RADIUS database.

Example: Configuring Static MAC Bypass of Authentication on an MX Series Router

Starting with Junos OS Release 14.2, to allow devices to access your LAN through 802.1X-configured interfaces without authentication, you can configure a static MAC bypass list on the MX Series router.The static MAC bypass list, also known as the *exclusion list*, specifies MAC addresses that are allowed on the router without a request to an authentication server.

You can use static MAC bypass of authentication to allow connection for devices that are not 802.1X-enabled, such as printers. If a host's MAC address is compared and matched against the static MAC address list, the nonresponsive host is authenticated and an interface opened for it.

This example describes how to configure static MAC bypass of authentication for two printers:

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.2 or later for MX240, MX480, or MX960 routers running in enhanced LAN mode.
- One router acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.

Before you connect the server to the router, be sure you have:

- Configured enhanced LAN mode on the router.
- Performed basic bridging and VLAN configuration on the router.
- Configured users on the RADIUS authentication server.

Overview and Topology

To permit printers access to the LAN, add them to the static MAC bypass list. The MAC addresses on this list are permitted access without authentication from the RADIUS server.

Consider an MX Series router that functions as an authenticator port. It is connected using the interface, ge-0/0/10, over the IP network to a RADIUS server. The router is also linked to a conference room using the interface, ge-0/0/1, to a printer using the interface, ge-0/0/20, to a hub using the interface, ge-0/0/8, and to two supplicants or clients over interfaces, ge-0/0/2 and ge-0/0/9 respectively.

The interfaces shown in Table 7 on page 92 will be configured for static MAC authentication.

Table 7: Components of the Static MAC Authentication Configuration Topology

Property	Settings
Router hardware	MX Series router
VLAN name	**default**
Connections to integrated printer/fax/copier machines (no PoE required)	**ge-0/0/19**, MAC address 00:04:0f:fd:ac:fe **ge-0/0/20**, MAC address 00:04:ae:cd:23:5f

The printer with the MAC address 00:04:0f:fd:ac:fe is connected to access interface **ge-0/0/19**. A second printer with the MAC address 00:04:ae:cd:23:5f is connected to access interface **ge-0/0/20**. Both printers will be added to the static list and bypass 802.1X authentication.

Configuration

CLI Quick Configuration

To quickly configure static MAC authentication, copy the following commands and paste them into the router terminal window:

```
[edit]
set protocols authentication-access-control static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f]
set protocols authentication-access-control interface all supplicant multiple
set protocols authentication-access-control authenticaton-profile-name profile1
```

Step-by-Step Procedure

Configure static MAC authentication:

1. Configure MAC addresses **00:04:0f:fd:ac:fe** and **00:04:ae:cd:23:5f** as static MAC addresses:

   ```
   [edit protocols]
   user@router# set authentication-access-control static [00:04:0f:fd:ac:fe
   00:04:ae:cd:23:5f]
   ```

2. Configure the 802.1X authentication method:

   ```
   [edit protocols]
   user@router# set authentication-access-control interface all supplicant multiple
   ```

3. Configure the authentication profile name (access profile name) to use for authentication:

   ```
   [edit protocols]
   ```



```
user@router# set authentication-access-control authentication-profile-name profile1
```

NOTE: Access profile configuration is required only for 802.1X clients, not for static MAC clients.

Results Display the results of the configuration:

```
user@router> show
interfaces {
  ge-0/0/19 {
    unit 0 {
      family bridge {
        vlan-id 10;
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      family bridge {
        vlan-id 10;
      }
    }
  }
}
protocols {
  authentication-access-control {
    authentication-profile-name profile1;
    static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f];
    interface {
      all {
        supplicant multiple;
      }
    }
  }
}
```

Verification

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

- Verifying Static MAC Bypass of Authentication on page 93

Verifying Static MAC Bypass of Authentication

Purpose Verify that the MAC address for both printers is configured and associated with the correct interfaces.

Action Use the operational mode command:

```
user@switch> show dot1x static-mac-address

MAC address          VLAN-Assignment              Interface
00:04:0f:fd:ac:fe    default                      ge-0/0/19.0
00:04:ae:cd:23:5f    default                      ge-0/0/20.0
```

Meaning The output field **MAC address** shows the MAC addresses of the two printers.

The output field **Interface** shows that the MAC address **00:04:0f:fd:ac:fe** can connect to the LAN through interface **ge-0/0/19.0** and that the MAC address **00:04:ae:cd:23:5f** can connect to the LAN through interface **ge-0/0/20.0**.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, to allow devices to access your LAN through 802.1X-configured interfaces without authentication, you can configure a static MAC bypass list on the MX Series router.

Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication on MX Series Routers

Starting with Junos OS Release 14.2, on MX Series routers, firewall filters that you apply to interfaces enabled for 802.1X or MAC RADIUS authentication are dynamically combined with the per-user policies sent to the switch from the RADIUS server.The switch uses internal logic to dynamically combine the interface firewall filter with the user policies from the RADIUS server and create an individualized policy for each of the multiple users or nonresponsive hosts that are authenticated on the interface.

This example describes how dynamic firewall filters are created for multiple supplicants on an 802.1X-enabled interface (the same principles shown in this example apply to interfaces enabled for MAC RADIUS authentication):

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.2 or later for MX Series routers
- One MX Series router

- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you apply firewall filters to an interface for use with multiple supplicants, be sure you have:

- Set up a connection between the router and the RADIUS server.
- Configured 802.1X authentication on the router, with the authentication mode for interface **ge-0/0/2** set to **multiple**.
- Configured users on the RADIUS authentication server.

Overview and Topology

When the 802.1X configuration on an interface is set to multiple supplicant mode, the system dynamically combines interface firewall filter with the user policies sent to the router from the RADIUS server during authentication and creates separate terms for each user. Because there are separate terms for each user authenticated on the interface, you can, as shown in this example, use counters to view the activities of individual users that are authenticated on the same interface.

When a new user (or a nonresponsive host) is authenticated on an interface, the system adds a term to the firewall filter associated with the interface, and the term (policy) for each user is associated with the MAC address of the user. The term for each user is based on the user-specific filters set on the RADIUS server and the filters configured on the interface. For example, as shown in Figure 5 on page 95, when User1 is authenticated by the MX Series router, the system creates the firewall filter **dynamic-filter-example**. When User2 is authenticated, another term is added to the firewall filter, and so on.

Figure 5: Conceptual Model: Dynamic Filter Updated for Each New User

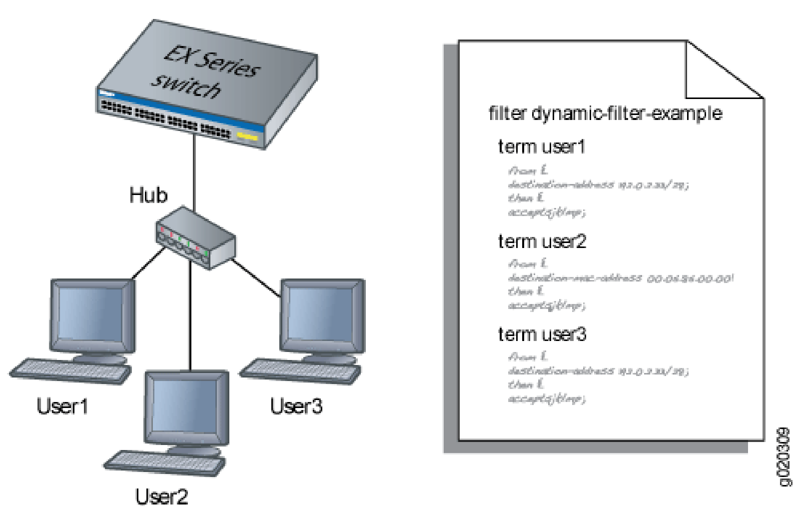

This is a conceptual model of the internal process—you cannot access or view the dynamic filter.

NOTE: If the firewall filter on the interface is modified after the user (or nonresponsive host) is authenticated, the modifications are not reflected in the dynamic filter unless the user is reauthenticated.

In this example, you configure a firewall filter to count the requests made by each endpoint authenticated on interface **ge-0/0/2** to the file server, which is located on subnet **192.0.2.16/28**, and set policer definitions to rate limit the traffic. Figure 6 on page 96 shows the network topology for this example.

Figure 6: Multiple Supplicants on an 802.1X-Enabled Interface Connecting to a File Server

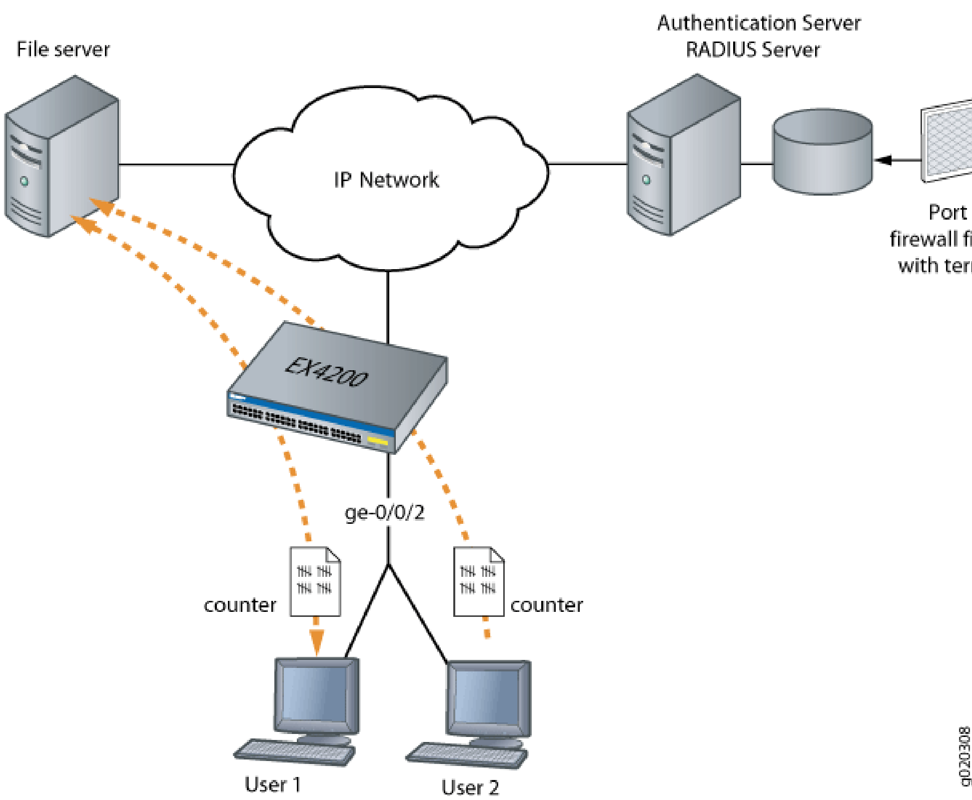

Configuration

To configure firewall filters for multiple supplicants on 802.1X-enabled interfaces:

- Configuring Firewall Filters on Interfaces with Multiple Supplicants on page 96

Configuring Firewall Filters on Interfaces with Multiple Supplicants

CLI Quick Configuration

To quickly configure firewall filters for multiple supplicants on an 802.1X-enabled interface copy the following commands and paste them into the router terminal window:

```
[edit]
set protocols authentication-access-control interface ge-0/0/2 supplicant multiple
set firewall family bridge filter filter1 term term1 from destination-address 192.0.2.16/28
```

set firewall policer p1 if-exceeding bandwidth-limit 1m
set firewall policer p1 if-exceeding burst-size-limit 1k
set firewall family bridge filter filter1 term term1 then count counter1
set firewall family bridge filter filter1 term term2 then policer p1

Step-by-Step Procedure

To configure firewall filters on an interface enabled for multiple supplicants:

1. Configure interface **ge-0/0/2** for multiple supplicant mode authentication:

   ```
   [edit protocols]
   user@router# set authentication-access-control interface ge-0/0/2 supplicant multiple
   ```

2. Set policer definition:

   ```
   user@router# show policer p1 |display set
   set firewall policer p1 if-exceeding bandwidth-limit 1m
   set firewall policer p1 if-exceeding burst-size-limit 1k
   set firewall policer p1 then discard
   ```

3. Configure a firewall filter to count packets from each user and a policer that limits the traffic rate. As each new user is authenticated on the multiple supplicant interface, this filter term will be included in the dynamically created term for the user:

   ```
   [edit firewall family bridge]
   user@router# set filter filter1 term term1 from destination-address 192.0.2.16/28
   user@router# set filter filter1 term term1 then count counter1
   user@router# set filter filter1 term term2 then policer p1
   ```

Results

Check the results of the configuration:

```
user@router> show configuration

firewall {
  family bridge {
    filter filter1 {
      term term1 {
        from {
          destination-address {
            192.0.2.16/28;
          }
        }
        then count counter1;
        term term2 {
        from {
          destination-address {
            192.0.2.16/28;
          }
        }
        then policer p1;
      }
    }
  }
```

```
    policer p1 {
        if-exceeding {
            bandwidth-limit 1m;
                burst-size-limit 1k;
            }
            then discard;
        }
}
protocols {
    authentication-access-control {
        interface ge-0/0/2 {
            supplicant multiple;
        }
    }
```

Verification

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

- Verifying Firewall Filters on Interfaces with Multiple Supplicants on page 98

Verifying Firewall Filters on Interfaces with Multiple Supplicants

Purpose Verify that firewall filters are functioning on the interface with multiple supplicants.

Action

1. Check the results with one user authenticated on the interface. In this case, the user is authenticated on **ge-0/0/2**:

```
user@router> show dot1x firewall
Filter: dot1x_ge-0/0/2
Counters
counter1_dot1x_ge-0/0/2_user1 100
```

2. When a second user, User2, is authenticated on the same interface, **ge-0/0/2**, you can verify that the filter includes the results for both of the users authenticated on the interface:

```
user@router> show dot1x firewall
Filter: dot1x-filter-ge-0/0/0
Counters
counter1_dot1x_ge-0/0/2_user1 100
counter1_dot1x_ge-0/0/2_user2 400
```

Meaning The results displayed by the **show dot1x firewall** command output reflect the dynamic filter created with the authentication of each new user. User1 accessed the file server located at the specified destination address 100 times, while User2 accessed the same file server 400 times.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, on MX Series routers, firewall filters that you apply to interfaces enabled for 802.1X or MAC RADIUS authentication are dynamically combined with the per-user policies sent to the switch from the RADIUS server.

CHAPTER 7

Configuring Aggregated Ethernet Interfaces for Increased Throughput and Link Redundancy

Aggregated Ethernet Interfaces Overview

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-flow load balancing.

NOTE: For information about configuring circuit cross-connects over aggregated Ethernet, see *Circuit and Translational Cross-Connects Overview*.

For information about mixed rates and mixed modes on an aggregated Ethernet bundle, see "Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles" on page 117.

Platform Support for Aggregated Ethernet Interfaces

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device and then associating a set of ports that have the same speed and are in full-duplex mode. The physical interfaces can be Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, or 10-Gigabit Ethernet IQ2 and IQ2-E. Generally, you cannot use a combination of these interfaces within the same aggregated link; however, you can combine Gigabit Ethernet and Gigabit Ethernet IQ interfaces in a single aggregated Ethernet bundle.

Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

TIP:

Going forward:

- Aggregated Ethernet link with member links of different modes will be referred as *10-Gigabit Ethernet mixed mode aggregated Ethernet link*.
- Aggregated Ethernet link with member links of different rates will be referred as *mixed rate aggregated Ethernet link*.
- These aggregated Ethernet links will generically be referred as *mixed aggregated Ethernet links*.

Table 8 on page 104 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 8: Platform Support Matrix for Mixed Aggregated Ethernet Bundles

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
10-Gigabit Ethernet LAN and WAN (WAN rate: OC192)	T640, T1600, T4000, and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (PF-24XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
		• Enhanced Scaling FPC3 (T640-FPC3-ES)	• 10-Gigabit Ethernet PIC with XENPAK (PC-1XGE-XENPAK)
		• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)
40-Gigabit Ethernet, 100-Gigabit Ethernet	T4000 and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 100-Gigabit Ethernet PIC with CFP (PF-1CGE-CFP)
	T640, T1600, T4000, and TX Matrix Plus routers	• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4) NOTE: This PIC is available packaged only in an assembly with the T1600-FPC4-ES FPC. • 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP)

Enhanced LAG Support on MX Series Routers

Starting in Junos OS Release 14.2, you can configure an enhanced link aggregation group (LAG) on MX Series routers. When you associate a physical interface with an aggregated Ethernet interface, the physical child links are also associated with the parent aggregated Ethernet interface to form a LAG.

In the absence of enhanced LAG support, one child next hop is created for each member link of an aggregated Ethernet interface for each VLAN interface. For example, an aggregate next hop for an aggregated Ethernet interface with 16 member links leads to the installation of 17 next hops per VLAN created. Thus the number of next hops supported on the routers with aggregated Ethernet interfaces is significantly reduced.

With the enhanced LAG support, when the [**edit chassis network-services enhanced-ip**] statement is configured, child next hops are not created for member links and, as a result, a higher number of next hops can be supported.

Note that the enhanced LAG feature is only supported when the router's network services is set to operate in the **enhanced-ip** mode. This feature is not supported when the router's network services is set to operate in the **enhanced-ethernet** mode.

Enhanced LAG Support on PTX Series Routers

Starting in Junos OS Release 18.1, Junos OS supports removal of child next hop usage for aggregated Ethernet Interfaces and clients on PTX Series routers with FPC3-PTX-U2 and FPC3-PTX-U3. Removal of child next hop usage helps reduce the memory and CPU resources required to support aggregated Ethernet Interfaces and improves the overall system performance and scaling numbers. This feature is enabled by default if the network services mode on the router is configured to **enhanced-mode**. You can disable this feature by using the **set chassis aggregated-devices disable-lag-enhanced**. You must reboot the router for the configuration to take effect.

Previously, each unicast next hop over aggregated Ethernet Interfaces resulted in creation of a number of children next hops as well. For an aggregated Ethernet Interface with 16 member links, addition of one unicast next hop over the aggregated Ethernet Interface results in installing total of 17 next hops. As a result, with aggregated Ethernet configuration, the number of next hops supported decreases in proportion to the number of aggregated Ethernet links.

NOTE: Child next hop optimizations are supported for aggregated Ethernet Interfaces, Interfaces that make use of aggregated Ethernet Interfaces, and for both unicast and multicast scenarios.

Configuration Guidelines for Aggregated Ethernet Interfaces

- Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, PICs, or MPCs.
- All Juniper routers support at least eight physical interfaces per aggregated Ethernet bundle. See maximum-links configuration page for platform specific limits.
- On M Series and T Series routers, you can create a maximum of 1024 logical interfaces on an aggregated Ethernet interface.
- Simple filters are not supported for interfaces in aggregated Ethernet bundles:
  - On M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
  - On MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

  For more information about simple filters, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.
- On the aggregated Ethernet bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific

capabilities, see “Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs” on page 529.

- Aggregated Ethernet interfaces can be either tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on MX Series routers support the configuration of **flexible-vlan-tagging** and **native-vlan-id** on dual-tagged frames, which consist of the following configuration statements:
  - inner-tag-protocol-id
  - inner-vlan-id
  - pop-pop
  - pop-swap
  - push-push
  - swap-push
  - swap-swap

  In all cases, you must set the number of aggregated Ethernet interfaces on the chassis. You can also set the link speed and the minimum links in a bundle.
- When configuring mixed aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers, consider the following:
  - A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
  - Link Aggregation Control Protocol (LACP), aggregated Ethernet link protection, and LACP link protection are supported only on mixed aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
  - Traffic distribution is based on the hash calculated on the egress packet header. Hash range is fairly distributed according to member links’ speed. This guarantees hash fairness but it does not guarantee fair traffic distribution depending on the rate of the egress streams.
  - Packets are dropped when the total throughput of the hash flow exiting a member link (or multiple hash flows exiting a single member link) exceeds the link speed of the member link. This can happen when egress member link changes because of a link failure and the hash flow switches to a member link of speed that is less than the total throughput of the hash flow.
  - Rate-based CoS components such as scheduler, shaper, and policer are not supported on mixed rate aggregated Ethernet links. However, the default CoS settings are supported by default on the mixed rate aggregated Ethernet links.
  - Load balancing is performed at the ingress Packet Forwarding Engine. Therefore, you must ensure that the egress traffic on the aggregated Ethernet link enters through the hardware platforms that support mixed aggregated Ethernet bundles. Table 8 on page 104 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

- Mixed aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed aggregated Ethernet load balancing is configured at egress.
- Load balancing of the egress traffic across the member links of a mixed rate aggregated Ethernet link is proportional to the rates of the member links.
- Egress multicast load balancing is not supported on mixed aggregated Ethernet interfaces.
- Changing the **edit interfaces ae*x* aggregated-ether-options link-speed** configuration of a mixed aggregated Ethernet link, which is configured on the supported interfaces of on T640, T1600, T4000, and TX Matrix Plus routers, leads to aggregated Ethernet link flapping.
- When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4), ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC with CFP, changing aggregated Ethernet link protection or LACP link protection configurations leads to aggregated Ethernet link flapping.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC with CFP, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP.
- The **show interfaces ae*x*** command displays the link speed of the aggregated Ethernet interface, which is the sum of the link speeds of all the active member links.

- Use the **show interfaces *aggregate-interface* extensive** and **show interfaces *aggregate.logical-interface*** commands to show the bandwidth of the aggregate. Also, the SNMP object identifier **ifSpeed/ifHighSpeed** shows the corresponding bandwidth on the aggregate logical interface if it is configured properly.

Release History Table

Release	Description
18.1	Starting in Junos OS Release 18.1, Junos OS supports removal of child next hop usage for aggregated Ethernet Interfaces and clients on PTX Series routers with FPC3-PTX-U2 and FPC3-PTX-U3.
14.2	Starting in Junos OS Release 14.2, you can configure an enhanced link aggregation group (LAG) on MX Series routers.
13.2	Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.

Related Documentation

- **inner-tag-protocol-id on page 1197**
- **inner-vlan-id on page 1198**
- **pop-pop on page 1307**
- **pop-swap on page 1308**
- **push-push on page 1328**
- **swap-push on page 1380**
- **swap-swap on page 1381**
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 124
- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 529
- *Ethernet Interfaces Feature Guide for Routing Devices*
- Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 117

Configuring an Aggregated Ethernet Interface

On Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on M Series and T Series routers, you can associate a physical interface with an aggregated Ethernet interface.

NOTE: On a Junos Fusion, you can include extended ports (physical interface on a satellite device that provides a connection to servers or endpoints) or local ports in link aggregation groups (LAGs) and MC-LAGs, but not both. For information on extended ports, see *Understanding Junos Fusion Ports* .

To configure an aggregated Ethernet interface:

1. Specify that you want to configure the link aggregation group interface.

   user@host# **edit interfaces** ***interface-name***

2. Configure the aggregated Ethernet interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | gigether-options) 802.3ad aex
```

You specify the interface instance number *x* to complete the link association; *x* can be from 0 through 127, for a total of 128 aggregated interfaces on M Series and T Series routers and can be from 1 through 480, for a total of 480 aggregated interfaces on MX Series routers. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. Aggregated interfaces are numbered from **ae0** through **ae4092**.

NOTE: On MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.

You must also include a statement defining **ae*x*** at the **[edit interfaces]** hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see "Ethernet Interfaces Overview" on page 3, and for a sample configuration, see "Example: Configuring Aggregated Ethernet Interfaces" on page 128.

NOTE: In general, aggregated Ethernet bundles support the features available on all supported interfaces that can become a member link within the bundle. As an exception, Gigabit Ethernet IQ features and some newer Gigabit Ethernet features are not supported in aggregated Ethernet bundles.

Gigabit Ethernet IQ and SFP interfaces can be member links, but IQ- and SFP-specific features are not supported on the aggregated Ethernet bundle even if all the member links individually support those features.

You need to configure the correct link speed for the aggregated Ethernet interface to eliminate any warning message.

NOTE: Before you commit an aggregated Ethernet configuration, ensure that link mode is not configured on any member interface of the aggregated Ethernet bundle; otherwise, the configuration commit check fails.

Related Documentation

- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 133
- Deleting an Aggregated Ethernet Interface on page 197
- Aggregated Ethernet Interfaces Overview on page 102

- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding Ethernet Link Aggregation on ACX Series Routers

Ethernet link aggregation is mechanism for increasing the bandwidth linearly and improving the resiliency of Ethernet links by bundling or combining multiple full-duplex same-speed point-to-point Ethernet links into a single virtual link. The virtual link interface is referred to as link aggregation group (LAG) or aggregated Ethernet (AE) interface. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

NOTE: ACX Series routers support connectivity fault management (CFM) on aggregated Ethernet interfaces with continuity check interval of 100 milliseconds or higher.

NOTE: ACX5048 and ACX5096 routers support connectivity fault management (CFM) on aggregated Ethernet interfaces with continuity check interval of 1 second or higher.

NOTE: The Ethernet options configurations for ACX5048 and ACX5096 routers differ compared to other ACX Series routers. For more information, see *Understanding Layer 2 Next Generation Mode on ACX Series Routers*.

On ACX Series routers, up to 128 AE interfaces can be created with each AE interface having up to 8 physical interfaces. AE interfaces can be created across PICs and fixed-ports on the chassis.

NOTE: On ACX5048 and ACX5096 routers, up to 64 AE interfaces can be created with each AE interface having up to 16 physical interfaces.

ACX Series routers do not support statistics for aggregated Ethernet interface. However, statistics can be retrieved for member interface.

To configure aggregated Ethernet interface:

1. Specify the number of aggregated Ethernet interfaces to be created:

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count number
   ```

2. Specify the minimum number of links for the aggregated Ethernet interface (ae*x*), that is, the defined bundle, to be labeled “up”:

NOTE: By default only one link must be up for the bundle to be labeled "up".

```
[edit interfaces]
user@host# set ae0 aggregated-ether-options minimum-links number (1 — 8)
```

3. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set ae0 aggregated-ether-options link-speed speed (10g | 1g | 100m)
```

4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set ge-1/0/0 gigether-options 802.3ad ae0
user@host# set ge-1/0/1 gigether-options 802.3ad ae0
```

5. Specify an interface family for the aggregated Ethernet bundle:

```
[edit interfaces]
user@host# set ae0 unit 0 family inet address ip-address
```

The above procedure creates an AE interface and they would be up and ready for running the services defined on AE logical interfaces.

AE interfaces can be VLAN-tagged or untagged. You can configure flexible-vlan-tagging, native-vlan-id, and dual-tagging on AE interfaces.

NOTE: Whenever there is a configuration change (AE interface to Gigabit Ethernet interfaces or vice versa), you need to remove the existing configuration, perform a commit, then add the new configuration and again commit the configuration.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

   This step changes the interface state to down and removes the configuration statements related to aex.

```
[edit]
user@host#delete interfaces aex
```

2. Delete the interface from the device count.

```
[edit]
user@host#delete chassis aggregated-devices ethernet device-count
```

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled.

Load Balancing

JUNOS load-balances traffic across member links in an AE bundle based on the Layer 3 information in the packet. You can globally configure what fields are used for load-balancing for inet and MPLS

On ACX Series Routers, the inet family knobs are available at PIC level. You can configure inet family Layer 3 and Layer 4 fields to be used for load-balancing. For bridge family, Layer 2, layer 3 and Layer 4 fields to be used for load-balancing.

ACX Series routers also support load balancing across the member links using Layer 2 source MAC addresses, destination MAC addresses, or both. This can be configured at the **[edit forwarding-options hash-key family multiservice]** hierarchy level. Layer 2 source MAC addresses and destination MAC addresses are used as hash-keys for load balancing.

```
[edit]
forwarding-options {
  hash-key {
    family multiservice {
      destination-mac;
      source-mac;
    }
  }
}
```

NOTE:

- For IP Layer 2 packets, only IP fields are used for load balancing across member links. Source MAC address and destination MAC address are not be used for load balancing.
- For non-IP Layer 2 packets, either Source MAC address or destination MAC address is used as hash-keys for load balancing.
- If you want to hash based on layer 2 fields, then you need to configure **multiservice**.
- If you want to hash based on layer 3 and layer 4 fields, then you need to configure **family (inet | inet6)**

LACP Monitoring

LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP is designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention

- Link monitoring to check whether both ends of the bundle are connected to the correct group

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

LACP monitoring can be either distributed or centralized. The default is distributed and it can be overriden by configuring the centralized knob under LACP protocols. LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

By default, LACP does not initiate a LACP PDU exchange. LACP packets can be configured to exchange LACP PDUs at a rate of 1 packet per second, or a slower rate of 1 packet for 30 seconds.

The LACP mode can be active or passive. If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is turned off on aggregated Ethernet interfaces. If LACP is configured, it is in passive mode by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode.

To enable LACP active mode, include the **lacp** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **active** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
  active;
}
```

NOTE: The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

To restore the default behavior, include the **lacp** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **passive** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
  passive;
}
```

Link Protection

Link protection can be configured on AE interfaces to provide 1:1 link resiliency using LACP. Primary and backup links can be configured within an AE bundle. The primary link is used for all transit traffic and host generated traffic. The backup link is used when the primary link fails.

Link protection is supported only when the AE bundles have no more than 2 member links, one primary and another backup. LACP works in revertive link-protection mode by default and can be configured to work in non-revertive mode.

NOTE: Link protection without LACP (static link protection on AE interfaces) is not supported on all ACX Series routers. Link protection works as expected with LACP configured on the AE bundle.

Configuring Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Configure the options for an aggregated Ethernet interface.

   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Configure the link protection mode.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **set link-protection**

Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert ae*x*** configuration command.

user@host# **delete interfaces ae*x* aggregated-ether-options link-protection**

Understanding the Algorithm Used to Hash LAG Bundle

ACX Series routers use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. The packet fields used by the hashing algorithm varies by the packet's EtherType and, in some instances, by the configuration on the router. The hashing algorithm recognizes the following EtherTypes:

- IPv4

- MPLS

Traffic that is not recognized as belonging to any of these EtherTypes is hashed based on the Layer 2 header. IP and MPLS traffic are also hashed based on the Layer 2 header when a user configures the hash mode as Layer 2 header.

You can configure some fields that are used by the hashing algorithm to make traffic forwarding decisions. You cannot, however, configure how certain values within a header are used by the hashing algorithm.

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.

Table 9 on page 115 describes the fields used for hashing by Layer 2 services. The table explains the default behavior and the configurable fields based on the type of traffic received on the Layer 2 service

Table 9: Hashing Behavior for Pseudowire (Layer 2 Circuit) and Bridging Services

Traffic Type	Default Hash Fields	Configurable Fields (Hash keys)
Layer 2	None	Source MAC Address Destination MAC Source MAC and Destination MAC
IP	Source IP and Destination IP	Source MAC Address Destination MAC Source MAC and Destination MAC
MPLS	MPLS label 1 and MPLS label 2	Source MAC Address Destination MAC Source MAC and Destination MAC

Table 10 on page 116 describes the fields used for hashing by Layer 3 services. The table explains the default behavior and the configurable fields based on the type of traffic received on the Layer 3 service

Table 10: Hashing Behavior for IP Services

Traffic Type	Default Hash Fields	Configurable Fields (Hash keys)
IP	Source IP and Destination IP	Layer 3 (Source IP and/or| destination IP) Layer 4 (UDP/TCP source port andr UDP/TCP destination port)

Related Documentation

- *CoS on ACX Series Universal Metro Routers Features Overview*
- *Controlling Network Access Using Traffic Policing Overview*
- *Firewall Filter Match Conditions and Actions on ACX Series Routers Overview*

Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a link aggregation group (LAG) or bundle. Link aggregation can be used for point-to-point connections. It balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

This topic describes how to configure aggregated Ethernet interfaces on PTX Series Packet Transport Routers.

On PTX Series Packet Transport Routers, aggregated Ethernet support includes the following features:

- A consistent interface type (**et *fpc/pic/port***) across all Ethernet interfaces.
- Ability to bundle multiple Ethernet interfaces
- Ability to bundle mixed-rate links on the same aggregated Ethernet interface
- Fault tolerance
- Load balancing between child links
- Advanced features including flexible VLAN tagging and Ethernet services encapsulation

Aggregated Ethernet interfaces can use interfaces from different FPCs or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
  aggregated-devices {
  ethernet {
    device-count 2;
  }
}
[edit interfaces]
et-0/0/0 {
```



```
gigether-options {
   802.3ad ae0;
   }
}
et-0/0/1 {
   gigether-options {
      802.3ad ae0;
   }
}
ae0 {
   vlan-tagging;
   unit 0 {
      vlan-id 100;
      family inet {
         address 200.200.1.2/24;
      }
   }
   unit 1 {
      vlan-id 101;
      family inet {
         address 200.200.2.2/24;
      }
   }
}
```

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 102
- Configuring Junos OS for Supporting Aggregated Devices on page 129
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 124

Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles

You can configure the member links of an aggregated Ethernet bundle with any combination of rates—also known as mixed rates—on T Series, MX Series, and PTX Series routers. The bandwidth that is provided by an aggregated Ethernet bundle can be utilized completely and efficiently when the links are configured with different rates.

You can configure mixed modes on T Series routers. In a mixed-mode configuration, the member links of an aggregated Ethernet bundle are configured in LAN mode as well as in WAN mode for 10-Gigabit Ethernet interfaces.

The following sections provide an overview about mixed rates and mixed modes on various platforms:

- Supported Platforms on page 118
- Aggregated Ethernet Bundle with Mixed Rates and Mixed Modes on T Series Routers on page 118
- Aggregated Ethernet Bundles with Mixed Rates on MX Series Routers on page 121

Supported Platforms

Table 11 on page 118 summarizes the first Junos OS release that supports aggregated Ethernet bundles on the various Juniper Networks routers and their components:

Table 11: Mixed Rates and Mixed Modes Support on Junos OS

Supported Platforms	Junos OS Release
Mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers	13.2
Mixed rates on MX240, MX480, MX960, MX2010, and MX2020 routers	14.1R2 14.2
Mixed rates on PTX Series Packet Transport Routers	14.2
Mixed rates on P2-10G-40G-QSFPP PIC and P2-100GE-OTN PIC on PTX5000 routers	15.1

Aggregated Ethernet Bundle with Mixed Rates and Mixed Modes on T Series Routers

The following sections explain mixed rates and mixed modes on T Series routers:

Understanding Mixed Rates and Mixed Modes

Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well.

  In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

Going forward:

- An aggregated Ethernet link with member links of different modes is referred to as *10-Gigabit Ethernet mixed-mode aggregated Ethernet link*.
- An aggregated Ethernet link with member links of different rates is referred to as *mixed-rate aggregated Ethernet link*.
- These aggregated Ethernet links will generically be referred to as *mixed aggregated Ethernet links*.

Platform Support Matrix for Mixed Aggregated Ethernet Bundles

Table 8 on page 104 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 12: Platform Support Matrix for Mixed Aggregated Ethernet Bundles

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
10-Gigabit Ethernet LAN and WAN (WAN rate: OC192)	T640, T1600, T4000, and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (PF-24XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
		• Enhanced Scaling FPC3 (T640-FPC3-ES)	• 10-Gigabit Ethernet PIC with XENPAK (PC-1XGE-XENPAK)
		• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)

Table 12: Platform Support Matrix for Mixed Aggregated Ethernet Bundles (continued)

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
40-Gigabit Ethernet, 100-Gigabit Ethernet	T4000 and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 100-Gigabit Ethernet PIC with CFP (PF-1CGE-CFP)
	T640, T1600, T4000, and TX Matrix Plus routers	• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4) NOTE: This PIC is available packaged only in an assembly with the T1600-FPC4-ES FPC. • 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP)

Guidelines to Follow When Configuring Aggregated Ethernet Bundles with Mixed Rates and Mixed Modes

In addition to the configuration guidelines for aggregated Ethernet interfaces in “Aggregated Ethernet Interfaces Overview” on page 102, you must consider the following as well when configuring mixed modes and mixed rates on aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers:

- A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
- Link Aggregation Control Protocol (LACP), aggregated Ethernet link protection, and LACP link protection are supported only on mixed aggregated Ethernet bundles configured on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Traffic distribution is based on the hash calculated on the egress packet header. Hash range is fairly distributed according to member links’ speed. This guarantees hash fairness but it does not guarantee fair traffic distribution depending on the rate of the egress streams.
- Packets are dropped when the total throughput of the hash flow exiting a member link (or multiple hash flows exiting a single member link) exceeds the link speed of the member link. This can happen when egress member link changes because of a link failure and the hash flow switches to a member link of speed that is less than the total throughput of the hash flow.
- Rate-based CoS components such as scheduler, shaper, and policer are not supported on mixed-rate aggregated Ethernet links. However, the default CoS settings are supported by default on the mixed-rate aggregated Ethernet links.
- Load balancing is performed at the ingress Packet Forwarding Engine. Therefore, you must ensure that the egress traffic on the aggregated Ethernet link enters through the hardware platforms that support mixed aggregated Ethernet bundles. Table 8 on page 104 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

- Mixed aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed aggregated Ethernet load balancing is configured at egress.
- Load balancing of the egress traffic across the member links of a mixed-rate aggregated Ethernet link is proportional to the rates of the member links.
- Egress multicast load balancing is not supported on mixed aggregated Ethernet interfaces.
- Changing the **[edit interfaces ae*x* aggregated-ether-options link-speed]** configuration of a mixed aggregated Ethernet link, which is configured on the supported interfaces of on T640, T1600, T4000, and TX Matrix Plus routers, leads to aggregated Ethernet link flapping.
- When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4), ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC with CFP, changing aggregated Ethernet link protection or LACP link protection configurations leads to aggregated Ethernet link flapping.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC with CFP, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP.
- The **show interfaces ae*x*** command displays the link speed of the aggregated Ethernet interface, which is the sum of the link speeds of all the active member links.

Aggregated Ethernet Bundles with Mixed Rates on MX Series Routers

The following sections explain mixed rates on aggregated Ethernet bundles:

Understanding Mixed Rates

Starting with Junos OS Release 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers, thereby enabling you to configure the member links with any combination of rates—10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet—on an aggregated Ethernet bundle.

You must configure the **mixed** statement explicitly at the **[edit interfaces *interface-name* aggregated-ether-options link-speed]** hierarchy level to:

- Enable the mixed-rate configuration on all the aggregated Ethernet bundles on the egress side of the Modular Port Concentrators (MPCs).

- Enable the router to detect child links of different speed values in the aggregated Ethernet bundle.

You can also configure the minimum bandwidth on an aggregated Ethernet bundle when you configure mixed rate on that aggregated Ethernet bundle.

NOTE: **The minimum-link statement cannot be configured with mixed rates.**

Mixed rates cannot be configured for aggregated Ethernet bundles on the egress side of the Dense Port concentrators (DPCs).

When you configure mixed rate on a homogeneous aggregated Ethernet bundle—where all the links in the bundle are of the same speed—the aggregated Ethernet bundle goes down and then comes up with the mixed-rate configuration.

Supported Features

The following features are supported on mixed-rate aggregated Ethernet bundles on MX Series routers:

- Sixty-four links in a bundle.
- Load balancing of traffic in proportion to the member-link speed.
- Non link-protect mode. For more information, see "Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection" on page 151.
- LACP for slow and fast interval for periodic transmission of LACP packets.
- Port-based network access control (NAC).
- Scheduler parameters for aggregated interface member links in a scaled manner with the **member-link-scheduler scale** statement at the **[edit class-of-service interfaces]** hierarchy level.
- Layer 3 features only.
- Configuration of following statements as percentages only for mixed rates at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level:
  - **buffer-size**
  - **excess-rate**
  - **shaping-rate**
  - **transmit-rate**
- Configuration of the following statements for mixed rates at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level:
  - **drop-profile-map**
  - **excess-priority**

  - **priority**
  - **transmit-rate (rate-limit | exact)**
- The **shared-bandwidth-policer** statement at the **[edit firewall policer *policer-name*]** hierarchy level.
- The **scheduler-maps** *map-name* statement at the **[edit class-of-service]** hierarchy level.
- Unicast load balancing, where the load balancing happens on ingress-only selectors.
- Mulitcast load balancing
- Make-before-break (MBB) for multicast LDP (MLDP) and fast reroute (FRR).
- Source class usage (SCU) and destination class usage (DCU) accounting.
- Families **inet**, **inet6**, and **mpls**.
- Enhanced IP network services.
- LDP tunneling and OAM link fault management (LFM).

The following features are not supported on mixed-rate aggregated Ethernet bundles on MX Series routers:

- Adaptive load balancing
- Hierarchical schedulers on aggregated Ethernet bundles and the scheduling on logical interfaces (per-unit scheduling).
- Shaping rate, where traffic shaping is achieved by specifying the amount of bandwidth to be allocated to a logical interface.
- The **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level.
- Ingress queuing.
- Options that are configured with nonpercentage values at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.
- The **member-link-scheduler replicate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level.
- Mixing LAN mode and WAN mode.
- Aggregated Ethernet link protection and link protection on a 1:1 model.
- LACP link protection.
- Layer 2 features.
- The **target-routing-instance (*routing-instance-name* | default)** statement at the **[edit access domain map *domain-map-name*]** hierarchy level.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers, thereby enabling you to configure the member links with any combination of rates—10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet—on an aggregated Ethernet bundle.
13.2	Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well.

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 102
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 124
- **minimum-bandwidth on page 1259**
- P2-10G-40G-QSFPP PIC Overview on page 409
- Understanding the P2-100GE-OTN PIC on page 500

Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles

The following sections explain how to configure mixed rates and mixed modes on various platforms:

Configuring Mixed Rates and Mixed Modes on an Aggregated Ethernet Bundle on T Series Routers

In releases before Junos OS Release 13.2, all interfaces that form an aggregated Ethernet bundle must have the same speed and must be in full-duplex mode. Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers. Following mixed rates and mixed modes are supported:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

NOTE:

- Member links of 50-Gigabit Ethernet can be configured using only the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of an aggregated Ethernet bundles in mixed rate or mixed mode on T640, T1600, T4000, and TX Matrix Plus routers:

1. Go to **[edit chassis]** hierarchy level.

   user@host# **edit chassis**

2. Configure the number of aggregated logical devices available to the router:

   [edit chassis]
   user@host# **set aggregated-devices ethernet device-count** *number*

3. Go to the **[edit interfaces]** hierarchy level.

   user@host# **top**
   user@host# **edit interfaces**

4. Configure the minimum number of links that are required for the aggregated Ethernet bundle to be labeled *up*:

   [edit interfaces]

user@host# **set ae*x* aggregated-ether-options minimum-links** *number*

NOTE: By default, only one link needs to be up for the bundle to be labeled *up*.

5. Configure the **link-speed** statement and specify the **mixed** option for the **link-speed** statement to indicate the mixed-rate and mixed-mode support for the aggregated Ethernet bundle configuration.

   [edit interfaces]
   user@host# **set ae*x* aggregated-ether-options link-speed mixed**

   NOTE: It is mandatory to configure the **mixed** option for aggregated Ethernet bundles for the PD-1CE-CFP-FPC4 PIC.

   On aggregated Ethernet bundles in MX Series routers, when the **mixed** statement at the **[edit interfaces ae*x* aggregated-ether-options link-speed]** hierarchy level is not configured, the mixed rate configuration is applied by default.

6. Configure the members links of the aggregated Ethernet bundle.

   [edit interfaces]
   user@host# **set** *interface-name* **gigether-options 802.3ad ae*x***

7. Configure an interface family and an IP address for the aggregated Ethernet bundle.

   [edit interfaces]
   user@host# **set ae*x* unit** *number* **family (inet | inet6 | mpls) address** *address*

8. Commit the configuration.

   [edit interfaces]
   user@host# **commit**

Configuring Mixed Rates on Aggregated Ethernet Bundles on MX Series Routers

Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers. You can now configure the member links with any combination of rates—10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet—on an aggregated Ethernet bundle.

To configure mixed rates on an aggregated Ethernet bundle on MX Series routers:

1. Go to the **[edit chassis]** hierarchy level.

   user@host# **edit chassis**

2. Configure the number of aggregated logical devices available to the router.

   [edit chassis]
   user@host# **set aggregated-devices ethernet device-count *number***

3. Go to the **[edit interfaces]** hierarchy level.

   user@host# **top**
   user@host# **edit interfaces**

4. Configure the **link-speed** statement and specify the **mixed** option for the **link-speed** statement to indicate the mixed-rate support for the aggregated Ethernet bundle configuration.

   [edit interfaces]
   user@host# **set ae*x* aggregated-ether-options link-speed mixed**

5. Configure the members links of the aggregated Ethernet bundle:

   [edit interfaces]
   user@host# **set *interface-name* gigether-options 802.3ad ae*x***

6. Configure an interface family for the aggregated Ethernet bundle as inet, inet6, or mpls:

   [edit interfaces]
   user@host# **set ae*x* unit *number* family (inet | inet6 | mpls)**

7. Configure the minimum bandwidth unit as **bps**, **gbps**, **kbps**, or **mbps** and the bandwidth value from 1 through 128,000.

   [edit interfaces]
   user@host# **set ae*x* aggregated-ether-options minimum-bandwidth bw-unit (bps | gbps | kbps | mbps)**
   user@host# **set ae*x* aggregated-ether-options minimum-bandwidth bw-value *value***

8. Commit the configuration.

   [edit interfaces]
   user@host# **commit**

Release History Table

Release	Description
14.1R1	Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.
13.2	Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers.
13.2	Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 102
- Configuring Aggregated Ethernet Link Speed on page 134
- Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers on page 116
- **link-speed on page 1229**
- **minimum-bandwidth on page 1259**
- Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 117

Example: Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count 15;
  }
}
[edit interfaces]
ge-1/3/0 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
```

```
            link-speed 1g;
            minimum-links 1;
        }
    }
    vlan-tagging;
    unit 0 {
        vlan-id 1;
        family inet {
            address 10.0.0.1/24;
        }
    }
    unit 1 {
        vlan-id 1024;
        family inet {
            address 10.0.0.2/24;
        }
    }
    unit 2 {
        vlan-id 1025;
        family inet {
            address 10.0.0.3/24;
        }
    }
    unit 3 {
        vlan-id 4094;
        family inet {
            address 10.0.0.4/24;
        }
    }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Configure 'link-speed' for Gigabit Ethernet based Aggregate Ethernet interface bundles

Configuring Junos OS for Supporting Aggregated Devices

Junos OS supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard.

Tasks for configuring aggregated devices are:

Configuring Virtual Links for Aggregated Devices

To define virtual links, you need to specify the associations between physical and logical devices within the **[edit interfaces]** hierarchy, and assign the correct number of logical devices by including the **device-count** statement at the **[edit chassis aggregated-devices ethernet]** and **[edit chassis aggregated-devices sonet]** hierarchy levels:

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count number;
  }
  sonet {
    device-count number;
  }
}
```

The aggregated interfaces are numbered from **ae0** through **ae4091**. The maximum number of aggregated interfaces supported by different routers is listed below:

- For PTX Series routers, you can configure a maximum of 128 aggregated interfaces.
- For M Series and T Series routers, you can configure a maximum of 128 aggregated interfaces (LAG bundles).
- In Junos release 14.2R2 and earlier, you can configure a maximum of 480 aggregated interfaces on MX Series routers.
- In Junos release 14.2R3 and later, you can configure a maximum of 1000 aggregated interfaces on MX240, MX480, and MX960 routers.
- In Junos release 14.2R3 and later, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.
- In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 480 aggregated interfaces on MX240, MX480, and MX960 routers.
- In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.

For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 64, numbered from **as0** through **as63**. In releases before Junos OS Release 13.2, the maximum was 16.

Configuring LACP Link Protection at the Chassis Level

Link Aggregation Control Protocol (LACP) is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled. LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP link protection enables you to force active and standby links within an aggregated Ethernet. You configure LACP link protection by using the **link-protection** and **system-priority** statements at either the chassis or interface level and by configuring port priority at the interface level using the **system-priority** statement. Configuring LACP parameters at the chassis level results in all aggregated Ethernet interfaces using the defined values unless overridden by the LACP configuration on a specific interface.

```
[edit chassis]
aggregated-devices {
  ethernet {
    lacp {
      link-protection {
        non-revertive;
      }
      system-priority priority;
    }
  }
}
```

NOTE: LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

See Also

Enabling LACP Link Protection

To enable LACP link protection for aggregated Ethernet interfaces on the chassis, use the **link-protection** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
link-protection {
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, after a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch, and the current link remains active.

BEST PRACTICE: (MX Series) In a highly scaled configuration over aggregated Ethernet, we recommend that you prevent the router from performing such a switch by including the **non-revertive** statement. Failure to do so may result in some traffic loss if a MIC on which a member interface is located reboots. Using the **non-revertive** statement for this purpose is not effective if both the primary and secondary interfaces are on the MIC that reboots.

CAUTION: **If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.**

Configuring System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the chassis, use the **system-priority** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
system-priority priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 through 65,535.

Configuring the Maximum Links Limit

To configure the maximum links limit, use the **maximum-links** statement at the **[edit chassis aggregated-devices]** hierarchy level:

```
[edit chassis aggregated-devices]
maximum-links maximum-links-limit;
```

Configuring PPM on Junos Fusion

If you use Junos Fusion with Junos OS Release 14.2R3, you need to ensure that link aggregation (and STP) work properly by configuring timers for the periodic packet management (PPM) daemons on the aggregation and satellite devices. We recommend using the following timer values:

```
[edit routing-options ppm]
  redistribution-timer 120;
  tcp-keepalive-interval 3000;
  tcp-keepalive-idle 3000;
```

Starting in Junos OS Release 14.2R4, the timer values that ensure proper link aggregation and STP functions are configured by default if you use Junos Fusion with Junos OS.

Release History Table

Release	Description
15.1F5	In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 480 aggregated interfaces on MX240, MX480, and MX960 routers.
15.1F5	In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.
14.2R4	Starting in Junos OS Release 14.2R4, the timer values that ensure proper link aggregation and STP functions are configured by default if you use Junos Fusion with Junos OS.
14.2R3	In Junos release 14.2R3 and later, you can configure a maximum of 1000 aggregated interfaces on MX240, MX480, and MX960 routers.
14.2R3	In Junos release 14.2R3 and later, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.
14.2R3	If you use Junos Fusion with Junos OS Release 14.2R3, you need to ensure that link aggregation (and STP) work properly by configuring timers for the periodic packet management (PPM) daemons on the aggregation and satellite devices.
14.2R2	In Junos release 14.2R2 and earlier, you can configure a maximum of 480 aggregated interfaces on MX Series routers.
13.2	For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 64, numbered from **as0** through **as63**.

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 108
- *Ethernet Interfaces Feature Guide for Routing Devices*
- Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers on page 116
- *Configuring Aggregated SONET/SDH Interfaces*

Configuring the Number of Aggregated Ethernet Interfaces on the Device

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them.

For M Series and T Series routers you can configure a maximum of 128 aggregated interfaces (LAG bundles). On MX Series routers running Junos release 14.2R2 and earlier, you can configure a maximum of 480 aggregated interfaces. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. For MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces. In all cases the aggregated interfaces are numbered from **ae0** through **ae4092**

NOTE: On a Junos Fusion Fabric, you can include extended ports (physical interface on a satellite device that provides a connection to servers or endpoints) or local ports in link aggregation groups (LAGs) and MC-LAGs, but not both. For information on extended ports, see *Understanding Junos Fusion Ports* .

For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 16, numbered from **as0** through **as15**.

1. Specify that you want to access the aggregated Ethernet configuration on the device.

   user@host# **edit chassis aggregated-devices ethernet**

2. Set the number of aggregated Ethernet interfaces.

   [edit chassis aggregated-devices ethernet]
   user@host# **set device-count** ***number***

You must also specify the constituent physical links by including the **802.3ad** statement at the **[edit interfaces *interface-name* fastether-options]** or **[edit interfaces *interface-name* gigether-options]** hierarchy level.

Release History Table

Release	Description
14.2R2	For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. For MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.
13.2	For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 16, numbered from **as0** through **as15**.

Related Documentation

- For information about physical links, see Configuring an Aggregated Ethernet Interface on page 108
- For a sample configuration, see Example: Configuring Aggregated Ethernet Interfaces on page 128
- *Ethernet Interfaces Feature Guide for Routing Devices*
- For information about configuring aggregated devices, see the *Junos OS Administration Library*.

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, there are exceptions.

Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers. For example, these mixes are supported:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.

Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of mixed rates and mixed modes on T640, T1600, T4000, TX Matrix Plus, and PTX routers, you need to configure the **mixed** option for the [**edit interfaces ae*x* aggregated-ether-options link-speed**] statement.

To set the required link speed:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces *interface-name*** aggregated-ether-options

2. Configure the link speed.

   [edit interfaces *interface-name* aggregated-ether-options]
   user@host# **set link-speed *speed***

   speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet interfaces on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.
- **50g**—Links are 50 Gbps.

Aggregated Ethernet links on T Series, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.
14.1	Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.
13.2	Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers.
13.2	Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.

Related Documentation

- **aggregated-ether-options on page 1079**
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 124
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces** ***interface-name*** **aggregated-ether-options**

2. Configure the minimum number of links.

   [edit interfaces *interface-name* aggregated-ether-options]
   user@host# **set minimum-links** ***number***

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links** ***number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links** ***number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links** ***number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On MX Series routers, when Link Aggregation Control Protocol (LACP) is enabled on a link aggregation group (LAG) interface along with minimum links configuration, the bundle is considered to be up when the following two conditions are met:

- The specified minimum number of links are up.
- The links are in *collecting distributing* state—that is, collecting and distributing states are merged together to form a combined state (coupled control) for the aggregated port. Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

Related Documentation

- aggregated-ether-options on page 1079
- minimum-links on page 1260

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces ae*x*]** hierarchy level:

```
[edit interfaces aex]
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

For more information about the **vlan-tagging** and **vlan-id** statements, see "802.1Q VLANs Overview" on page 244.

Related Documentation

- vlan-id on page 1432
- vlan-tagging on page 1445

Configuring Untagged Aggregated Ethernet Interfaces

Packet tagging provides a logical way to differentiate traffic on ports which support multiple virtual local area network (VLAN). While you must configure aggregated Ethernet interfaces to receive tagged traffic, you must also configure aggregated Ethernet interfaces that can receive untagged traffic.

To configure an aggregated Ethernet interface as untagged, remove the **vlan-tagging** statement at the **[edit interfaces ae*x*]** hierarchy level and remove the **vlan-id** statement from the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

NOTE: **You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.**

Table 13 on page 139 lists untagged aggregated Ethernet and LACP support by PIC and router.

Table 13: Untagged Aggregated Ethernet and LACP Support by PIC and Platform

PIC Type	M Series	LACP	T Series	LACP
4-port Fast Ethernet PIC Type 1	Yes	Yes	Yes	Yes
1-port Gigabit Ethernet PIC Type 1	Yes	Yes	Yes	Yes
2-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
4-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
1-port 10-Gigabit Ethernet M160	Yes	Yes	NA	NA
10-port Gigabit Ethernet PIC Type 3	Yes (M120, M320)	Yes	Yes	Yes
1-port 10-Gigabit Ethernet PIC Type 3	N/A	NA	Yes	Yes
8-port Gigabit Ethernet PIC Type 3	Yes	Yes	Yes	Yes

The 8–port Fast Ethernet PIC does not support untagged aggregated Ethernet or LACP.

Syslog messages are logged if you try to configure an untagged aggregated Ethernet interface using an unsupported PIC type.

Related Documentation

- Configuring Tagged Aggregated Ethernet Interfaces on page 138
- Configuring LACP for Aggregated Ethernet Interfaces on page 140.

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring LACP for Aggregated Ethernet Interfaces

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled.

NOTE: Starting with Junos OS Release 14.1, you can configure aggregated Ethernet interfaces with LACP on logical systems within an MX Series router.

For Multichassis Link Aggregation (MC-LAG), you must specify the **system-id** and **admin key**. MC-LAG peers use the same **system-id** while sending the LACP messages. The **system-id** can be configured on the MC-LAG network device and synchronized between peers for validation.

LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

The LACP mode can be active or passive. If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is turned off on aggregated Ethernet interfaces. If LACP is configured, it is in passive mode by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode.

NOTE: LACP can link together multiple different physical interfaces, but only features that are supported across all of the linked devices will be supported in the resulting link aggregation group (LAG) bundle. For example, different PICs can support a different number of forwarding classes. If you use link aggregation to link together the ports of a PIC that supports up to 16 forwarding classes with a PIC that supports up to 8 forwarding classes, the resulting LAG bundle will only support up to 8 forwarding classes. Similarly, linking together a PIC that supports WRED with a PIC that does not support it will result in a LAG bundle that does not support WRED.

To enable LACP active mode, include the **lacp** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **active** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
  active;
}
```

NOTE: The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

If you restart the Link Aggregation Control Protocol (LACP) process consecutively without adequate sleep or pause time between the successive restarts, the LACP links might flap. To avoid this problem, you must restart the LACP process only after the refresh time of the periodic packet management (PPM) process is completed.

To restore the default behavior, include the **lacp** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **passive** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
  passive;
}
```

Starting with Junos OS Release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.

To override the IEEE 802.3ad standard and facilitate subsecond failover, include the **fast-failover** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you use the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

For more information, see the following sections:

Configuring the LACP Interval

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level:

[edit interfaces *interface-name* aggregated-ether-options lacp]
periodic *interval*;

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

NOTE: Starting with Junos OS Release 11.4, source address filtering does not work when LACP is enabled., This behavior is not applicable to T Series routers and PTX Series Packet Transport Routers. For more information about source address filtering, see "Configuring MAC Address Filtering for Ethernet Interfaces" on page 14.

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see "Configuring Untagged Aggregated Ethernet Interfaces" on page 138.

For M Series Multiservice Edge Routers with enhanced Flexible PIC Concentrators (FPCs) and T Series routers, LACP over VLAN-tagged aggregated Ethernet interfaces is supported. For 8-port, 12-port, and 48-port Fast Ethernet PICs, LACP over VLAN-tagged interfaces is not supported.

LACP **Fast Periodic**, which is achieved by configuring fast (every second) intervals for periodic transmission of LACP packets, is supported with graceful Routing Engine switchover (GRES) on MX Series routers only.

Configuring LACP Link Protection

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the **link-protection** and **system-priority** statements. Configuring values at this level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

NOTE:

- When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.
- LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

Enabling LACP Link Protection

To enable LACP link protection for an aggregated Ethernet interface, use the **link-protection** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable;
  revertive;
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, once a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch and the current link remains active.

If LACP link protection is configured to be nonrevertive at the global (**[edit chassis]** hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.

CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

We strongly recommend that you use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces of MX Series routers to any other vendor device. Otherwise, the vendor device (say a Layer 2 switch, or a router) will not be able to manage the traffic coming from the two link aggregated Ethernet bundle. As a result, you might observe the vendor device sending back the traffic to the backup member link of the aggregated Ethernet interface.

Currently, MX-MPC2-3D, MX-MPC2-3D-Q, MX-MPC2-3D-EQ, MX-MPC1-3D, MX-MPC1-3D-Q, and MPC-3D-16XGE-SFPP do not drop traffic coming back to the backup link, whereas DPCE-R-Q-20GE-2XGE, DPCE-R-Q-20GE-SFP, DPCE-R-Q-40GE-SFP, DPCE-R-Q-4XGE-XFP, DPCE-X-Q-40GE-SFP, and DPCE-X-Q-4XGE-XFP drop traffic coming to the backup link.

Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 to 65,535.

Configuring LACP System Identifier

To configure the LACP system identifier for aggregated Ethernet interfaces, use the **system-id** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-id system-id;
```

Starting with Junos OS Release 13.3, youmust not configure the LACP system identifier by using the **system-id *system-id*** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level to be all zeros (00:00:00:00:00:00). If you attempt to commit a configuration with the system identifier to be all zeros, an error occurs during the commit operation.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Configuring LACP administrative Key

To configure an administrative key for LACP, include the **admin-key *number*** statement at the **[edit interfaces ae*x* aggregated-ether-options** lacp**]** hierarchy level:

```
[edit interfaces ae x aggregated-ether-options-lacp]
admin-key number;
```

NOTE: **You must configure MC-LAG to configure the admin-key statement. For more information about MC-LAG, see *Configuring Multichassis Link Aggregation on MX Series Routers*.**

Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the **port-priority** statement at the **[edit interfaces *interface-name* gigether-options 802.3ad ae*X* lacp]** or **[edit interfaces *interface-name* fastether-options 802.3ad ae*X* lacp]** hierarchy levels:

```
[edit interfaces interface-name gigether-options 802.3ad aeX lacp]
port-priority priority;
```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and is assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.

NOTE: **Starting with Junos OS Release 9.3, port aggregation selection (discussed previously) is performed for the active link when LACP link protection is enabled. Without LACP link protection, port priority is not used in port aggregation selection.**

Configuring LACP Hold-Up Timer to Prevent Link Flapping on LAG Interfaces

On link aggregation group (LAG) interfaces, when a member (child) link goes down, its state changes from current to expired. This link might flap from the current state to the expired state and back to current state when it receives intermittent LACP protocol data units (PDUs) and keepalive timeouts. Such flapping can adversely affect the traffic on the link.

To prevent excessive flapping of a LAG child link, you can configure a hold-up timer on the LAG interface that is applicable to all member links on that particular interface. To hold up, in networking terms, means to prevent the transitioning of an interface from down to up for a specified time interval.

When configured, the hold-up timer is triggered when an LACP state machine tries to move to the current state from the expired or default state when it receives an LACP PDU. The hold-up timer is triggered only if the LACP state machine had acquired the current state at least once earlier. The timer is not triggered if LACP attempts to transition to the current state for the first time. LACP monitors the PDUs received on the child link but prevents the link from transitioning to current state. If no flapping is observed when the link receives the PDUs, the hold-up timer expiries and triggers the member link to transition back to the current state. This transition is triggered as soon as the hold-up timer expires and not necessarily when the link receives a PDU.

To configure LACP hold-up timer for LAG interface, use the **hold-time up** statement at the **[edit interfaces aex aggregated-ether-options lacp]** hierarchy level.

NOTE:

- The hold-up timer keeps running even when the interface that receives the LACP PDU moves to the port disable state. The timer is then restarted if, before the timer expires, the interface comes up again and receives an LACP PDU from its neighbor. This ensures that the timer is maintained even during a quick physical port flap.
- When the following events occur, a hold-up timer is not triggered until the member link acquires the current state after the event:
  - LACP daemon restart
  - Deactivation and reactivation of child or aggregated Ethernet interface
  - Deletion and reconfiguration of child or aggregated Ethernet interface
  - System reboot
  - Routing Engine switchover

Tracing LACP Operations

To trace the operations of the LACP process, include the **traceoptions** statement at the **[edit protocols lacp]** hierarchy level:

```
[edit protocols lacp]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```


Sample Configuration for Configuring Aggregated Ethernet LACP on Tagged and Untagged Interfaces

Following configurations are examples of configuring aggregated Ethernet LACP on VLAN-tagged and untagged interfaces:

LACP with VLAN-Tagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}
```

LACP with Untagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
```

}

Release History Table

Release	Description
14.1	Starting with Junos OS Release 14.1, you can configure aggregated Ethernet interfaces with LACP on logical systems within an MX Series router.
13.3	Starting with Junos OS Release 13.3, youmust not configure the LACP system identifier by using the **system-id** ***system-id*** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level to be all zeros (00:00:00:00:00:00).
12.2	Starting with Junos OS Release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.
11.4	Starting with Junos OS Release 11.4, source address filtering does not work when LACP is enabled.
9.3	Starting with Junos OS Release 9.3, port aggregation selection (discussed previously) is performed for the active link when LACP link protection is enabled.

Related Documentation

- *Junos OS Administration Library*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Aggregated Ethernet Link Protection

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

NOTE: **Link protection is not supported on MX80.**

Configuring Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Configure the link protection mode.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **set link-protection**

See Also

- Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 149
- Example: Configuring Aggregated Ethernet Link Protection on page 150
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* primary**

2. Configure the backup logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* backup**

See Also

- 802.3ad on page 1063
- *Ethernet Interfaces Feature Guide for Routing Devices*

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

user@host> **request interface revert ae*x***

See Also

- Configuring Link Protection for Aggregated Ethernet Interfaces on page 149
- *Ethernet Interfaces Feature Guide for Routing Devices*

Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert ae*x*** configuration command.

user@host# **delete interfaces ae*x* aggregated-ether-options link-protection**

See Also

- Configuring Link Protection for Aggregated Ethernet Interfaces on page 149
- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring Aggregated Ethernet Link Protection

The following configuration enables link protection on the **ae0** interface, and specifies the **ge-1/0/0** interface as the primary link and **ge-1/0/1** as the secondary link.

```
[edit interfaces]
ae0 {
  aggregated-ether-options {
    link-protection;
  }
}
[edit interfaces]
ge-1/0/0 {
  gigether-options {
    802.3ad ae0 primary;
  }
}
[edit interfaces]
ge-1/0/1 {
  gigether-options {
    802.3ad ae0 backup;
  }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Shared Scheduling on Aggregated Ethernet Interfaces

You can configure shared scheduling on aggregated Ethernet interfaces in link protection mode on Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Ethernet Enhanced IQ2 (IQ2E) PICs on M320 routers.

To configure shared scheduling on aggregated Ethernet interfaces:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Configure the link protection mode.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **set link-protection**

3. Configure shared scheduling.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **top**
   [edit]
   user@host# **edit interfaces ae*x* shared-scheduler**

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 148

Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection

On aggregated Ethernet interfaces, you can configure scheduler in non–link-protect mode on the following platforms:

- MX-Series
- M120 and M320 with IQ2 PIC
- T-series platforms (T620 and T320) with IQ2 PIC

The scheduler functions supported are:

- Per unit scheduler
- Hierarchical scheduler
- Shaping at the physical interface

To configure the hierarchical scheduler on aggregated Ethernet interfaces in the non link-protect mode, include the **hierarchical-scheduler** statement at the **[edit interfaces ae*X*]** hierarchy level:

[edit interfaces ae*X* hierarchical-scheduler]

Prior to Junos OS Release 9.6, the hierarchical scheduler mode on these models required the **aggregated-ether-options** statement **link-protection** option. If a **link-protection** option is not specified, the scheduler is configured in non-link-protect mode.

To specify the member link bandwidth derivation based on the equal division model **(scale)** or the replication model **(replicate)** on aggregated Ethernet interfaces, include the **member-link-scheduler (scale | replicate)** option at the **[edit class-of-service interfaces ae*X*]** hierarchy level. The default setting is **scale**.

```
[edit class-of-service interfaces aeX member-link-scheduler (scale | replicate)]
```

NOTE: In link-protect mode, only one link is active at a time and the other link acts as the backup link, whereas in a non link-protect mode, all the links of the aggregate bundle are active at the same time. There is no backup link. If a link goes down or a new link is added to the bundle, traffic redistribution occurs.

Related Documentation

- *Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- For more information on the hierarchical scheduler (CoS), see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview

MX Series routers with Aggregated Ethernet PICs support symmetrical load balancing on an 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through deep packet inspection (DPI) devices over an LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Without symmetrical load balancing on an 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. By using this feature, a given flow of traffic (duplex) is ensured for the same devices in both directions.

Symmetrical load balancing on an 802.3ad LAG utilizes a mechanism of interchanging the source and destination addresses for a hash computation of fields, such as source address and destination address. The result of a hash computed on these fields is used to choose the link of the LAG. The hash-computation for the forward and reverse flow must be identical. This is achieved by swapping source fields with destination fields for the reverse flow. The swapped operation is referred to as *complement hash computation* or **symmetric-hash complement** and the regular (or unswapped) operation as *symmetric-hash computation* or **symmetric-hash**. The swappable fields are MAC address, IP address, and port.

Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

You can specify whether symmetric hash or complement hash is done for load-balancing traffic. To configure symmetric hash, use the **symmetric-hash** statement at the **[edit forwarding-options hash-key family inet]** hierarchy level. To configure symmetric hash complement, use the **symmetric-hash complement** statement and option at the **[edit forwarding-options hash-key family inet]** hierarchy level.

These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the **symmetric-hash** or **symmetric-hash complement** statement at the **[edit chassis hash-key family inet]** and **[edit chassis hash-key family multiservice]** hierarchy levels.

Consider the example in Figure 7 on page 153.

Figure 7: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

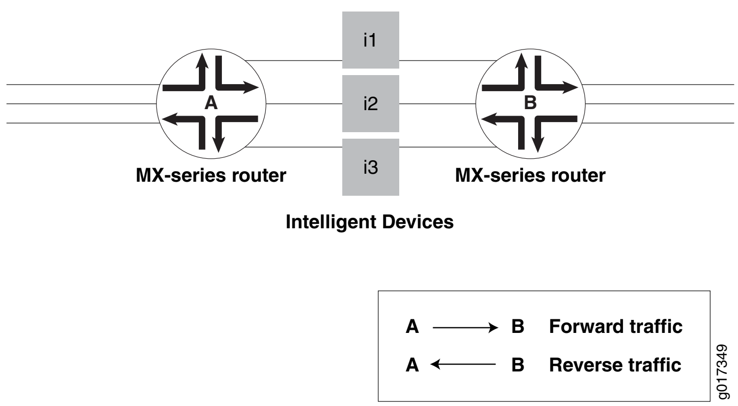

Router A is configured with symmetric hash and Router B is configured with symmetric hash complement. Thus, for a given flow *fx*, post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow *fx* is from Router B to Router A through the same i2 device as its hashing (done after swapping source and destination fields) and returns the same link index; since it is performed on the interchanged source and destination addresses.

However, the link chosen may or may not correspond to what was attached to the DPI. In other words, the hashing result should point to the same links that are connected, so that the traffic flows through the same DPI devices in both directions. To make sure this happens, you need to also configure the counterpart ports (ports that are connected to same DPI-iN) with the identical link index. This is done when configuring a child-link into the LAG bundle. This ensures that the link chosen for a given hash result is always the same on either router.

Note that any two links connected to each other should have the same link index and these link indices must be unique in a given bundle.

NOTE:
The following restrictions apply when configuring symmetric load balancing on an 802.3ad LAG on MX Series routers:

- The Packet Forwarding Engine (PFE) can be configured to hash the traffic in either symmetric or complement mode. A single PFE complex cannot work simultaneously in both operational modes and such a configuration can yield undesirable results.
- The per-PFE setting overrides the chassis-wide setting only for the family configured. For the other families, the PFE complex still inherits the chassis-wide setting (when configured) or the default setting.
- This feature supports VPLS, INET, and bridged traffic only.
- This feature cannot work in tandem with the **per-flow-hash-seed load-balancing** option. It requires that all the PFE complexes configured in complementary fashion share the same seed. A change in the seed between two counterpart PFE complexes may yield undesired results.

For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library*.

Example Configuration Statements

To configure 802.3ad LAG parameters at the bundle level:

```
[edit interfaces]
g(x)e-fpc/pic/port {
  gigether-options {
    802.3ad {
      bundle;
      link-index number;
    }
  }
}
```

where the **link-index *number*** ranges from 0 through 15.

You can check the link index configured above using the **show interfaces** command:

```
[edit forwarding-options hash-key]
family inet {
  layer-3;
  layer-4;
  symmetric-hash {
    [complement;]
  }
}
family multiservice {
  source-mac;
  destination-mac;
```



```
    payload {
    ip {
      layer-3 {
        source-ip-only | destination-ip-only;
      }
      layer-4;
    }
    }
    symmetric-hash {
      [complement;]
    }
}
```

For load-balancing Layer 2 traffic based on Layer 3 fields, you can configure 802.3ad LAG parameters at a per PIC level. These configuration options are available under the chassis hierarchy as follows:

```
[edit chassis]
fpc X {
  pic Y {
    .
    .
    .
    hash-key {
      family inet {
        layer-3;
        layer-4;
        symmetric-hash {
          [complement;]
        }
      }
      family multiservice {
        source-mac;
        destination-mac;
        payload {
          ip {
            layer-3 {
              source-ip-only | destination-ip-only;
            }
            layer-4;
          }
        }
        symmetric-hash {
          [complement;]
        }
      }
    }
    .
    .
    .
  }
}
```


Configuring Symmetrical Load Balancing on Trio-Based MPCs

With some configuration differences, symmetrical load-balancing over an 802.3ad link aggregation group is supported on MX Series routers with Trio-based MPCs.

To achieve symmetrical load-balancing on Trio-Based MPCs, the following needs to be done:

- Compute a Symmetrical Hash

  Both routers must compute the same hash value from the flow in the forward and reverse directions.On Trio-based platforms, the calculated hash value is independent of the direction of the flow, and hence is always symmetric in nature. For this reason, no specific configuration is needed to compute a symmetric hash value on Trio-based platforms.

  However, it should be noted that the fields used to configure the hash should have identical include and exclude settings on both ends of the LAG.

- Configure Link Indexes

  To allow both routers to choose the same link using the same hash value, the links within the LAG must be configured with the same link index on both routers. This can be achieved with the **link-index** statement.

- Enable Symmetric Load Balancing

  To configure symmetric load balancing on Trio-based MPCs, include the **symmetric** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy level. This statement is applicable to Trio-based platforms only.

  The **symmetric** statement can be used with any protocol family and enables symmetric load-balancing for all aggregated Ethernet bundles on the router. The statement needs to be enabled at both ends of the LAG. This statement is disabled by default.

- Achieve Symmetry for Bridged and Routed Traffic

  In some deployments, the LAG bundle on which symmetry is desired is traversed by Layer 2 bridged traffic in the upstream direction and by IPv4 routed traffic in the downstream direction. In such cases, the computed hash is different in each direction because the Ethernet MAC addresses are taken into account for bridged packets. To overcome this, you can exclude source and destination MAC addresses from the enhanced-hash-key computation.

  To exclude source and destination MAC addresses from the enhanced-hash-key computation, include the **no-mac-addresses** statement at the **[edit forwarding-options enhanced-hash-key family multiservice]** hierarchy level. This statement is disabled by default.

When symmetrical load balancing is enabled on Trio-based MPCs, keep in mind the following caveats:

- Traffic polarization is a phenomenon that occurs when using topologies that distribute traffic by using hashing of the same type. When routers are cascaded, traffic polarization can occur, and this can lead to unequal traffic distribution.

Traffic polarization occurs when LAGs are configured on cascaded routers. For example, in Figure 8 on page 157, if a certain flow uses Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, the flow also uses Link 1 of the aggregated Ethernet bundle between Device R2 and Device R3.

Figure 8: Traffic Polarization on Cascaded Routers When Symmetrical Load Balancing in Enabled on Trio-based MPCs

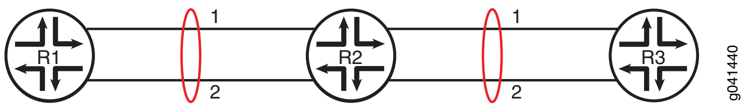

This is unlike having a random link selection algorithm, where a flow might use Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, and Link 2 of the aggregated Ethernet bundle between Device R2 and Device R3.

- Symmetric load balancing is not applicable to per-prefix load-balancing where the hash is computed based on the route prefix.
- Symmetric load balancing is not applicable to MPLS or VPLS traffic, because in these scenarios the labels are not the same in both directions.

Example Configurations

- Example Configurations of Chassis Wide Settings on page 157
- Example Configurations of Per–Packet-Forwarding-Engine Settings on page 158

Example Configurations of Chassis Wide Settings

Router A

```
user@host> show configuration forwarding-options hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric hash;
}
```

Router B

```
user@host> show configuration forwarding-options hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric-hash {
    complement;
  }
}
```

Example Configurations of Per–Packet-Forwarding-Engine Settings

Router A

```
user@host> show configuration chassis fpc 2 pic 2 hash-key
family multiservice {
   payload {
      ip {
         layer-3;
      }
   }
   symmetric hash;
}
```

Router B

```
user@host> show configuration chassis fpc 2 pic 3 hash-key
family multiservice {
   payload {
      ip {
         layer-3;
      }
   }
   symmetric-hash {
      complement;
   }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library*.

Understanding Aggregated Ethernet Load Balancing

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

  By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

  ECMP load balancing can be:

  - Across BGP paths (BGP multipath)
  - Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

  Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

  To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  NOTE: Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers and QFX10000 switches only.

To configure the tolerance value as a percentage, include the **tolerance** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the **pps** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the **scan-interval** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

NOTE: The **pps** and **scan-interval** optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-Packet Random Spray Load Balancing

  When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  NOTE: The Per-Packet option for load balancing is not supported on PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the **show interfaces ae*x* aggregated-ether-options load-balance** command.

Related Documentation

- Example: Configuring Aggregated Ethernet Load Balancing on page 163

Example: Configuring Aggregated Ethernet Load Balancing

- Understanding Aggregated Ethernet Load Balancing on page 161
- Example: Configuring Aggregated Ethernet Load Balancing on page 163

Understanding Aggregated Ethernet Load Balancing

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

  Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

  By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

  ECMP load balancing can be:

  - Across BGP paths (BGP multipath)
  - Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

  Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

  To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  NOTE: **Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers and QFX10000 switches only.**

  To configure the tolerance value as a percentage, include the **tolerance** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the **pps** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the **scan-interval** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  NOTE: **The pps and scan-interval optional keywords are supported on PTX Series Packet Transport Routers only.**

- Per-Packet Random Spray Load Balancing

  When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  NOTE: **The Per-Packet option for load balancing is not supported on PTX Series Packet Transport Routers.**

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the **show interfaces ae*x* aggregated-ether-options load-balance** command.

See Also

- Example: Configuring Aggregated Ethernet Load Balancing on page 163

Example: Configuring Aggregated Ethernet Load Balancing

This example shows how to configure aggregated Ethernet load balancing.

Requirements

This example uses the following hardware and software components:

- Three MX Series routers with MIC and MPC interfaces or three PTX Series Packet Transport Routers with PIC and FPC interfaces
- Junos OS Release 13.3 or later running on all devices

Overview

Load balancing is required on the forwarding plane when there are multiple paths or interfaces available to the next hop router, and it is best if the incoming traffic is load balanced across all available paths for better link utilization.

Aggregated Ethernet bundle is a typical application that uses load balancing to balance traffic flows across the member links of the bundle (IEEE 802.3ad).

Starting with Junos OS Release 13.3, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on MICs or MPCs of MX Series routers. Starting with Junos OS Release 14.1, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on PICs or FPCs of PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are:

- Adaptive—Adaptive load balancing is used in scenarios where flow-based hashing is not sufficient to achieve a uniform load distribution. This load-balancing solution implements a real-time feedback and control mechanism to monitor and manage imbalances in network load.

  The adaptive load-balancing solution corrects the traffic flow imbalance by modifying the selector entries, and periodically scanning the link utilization on each member link of the AE bundle to detect any deviations. When a deviation is detected, an adjustment event is triggered and fewer flows are mapped to the affected member link. As a result, the offered bandwidth of that member link goes down. This causes a continuous feedback loop, which over a period of time ensures that the same amount of byte rate is offered to all the member links, thus providing efficient traffic distribution across each member link in the AE bundle.

To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

NOTE: Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers only.

The **pps** option enables load balancing based on the packets-per-second rate. The default setting is bits-per-second load balancing.

The **scan-interval** value configures the length of time for scanning as a multiple of 30 seconds.

The **tolerance** value is the limit to the variance in the packet traffic flow to the aggregated Ethernet links in the bundle. You can specify a maximum of 100-percent variance. When the tolerance attribute is not configured, a default value of 20 percent is enabled for adaptive load balancing. A smaller tolerance value balances better bandwidth, but takes a longer convergence time.

NOTE: The pps and scan-interval optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-packet random spray—When the adaptive load-balancing solution fails, per-packet random spray acts as a last resort. The per-packet random spray load-balancing solution helps to address traffic imbalance by randomly spraying the packets to the aggregate next hops. This ensures that all the member links of the AE bundle are equally loaded, resulting in packet reordering.

  In addition, per-packet random spray identifies the ingress Packet Forwarding Engine that caused the traffic imbalance and eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  NOTE: The Per-Packet option for load balancing is not supported on the PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being implemented by issuing the **show interfaces ae*x* aggregated-ether-options load-balance** command.

Topology

In this topology, two aggregated Ethernet bundles - ae0 and ae1 - are configured on the links between the R2 and R3 routers.

Figure 9: Aggregated Ethernet Load Balancing

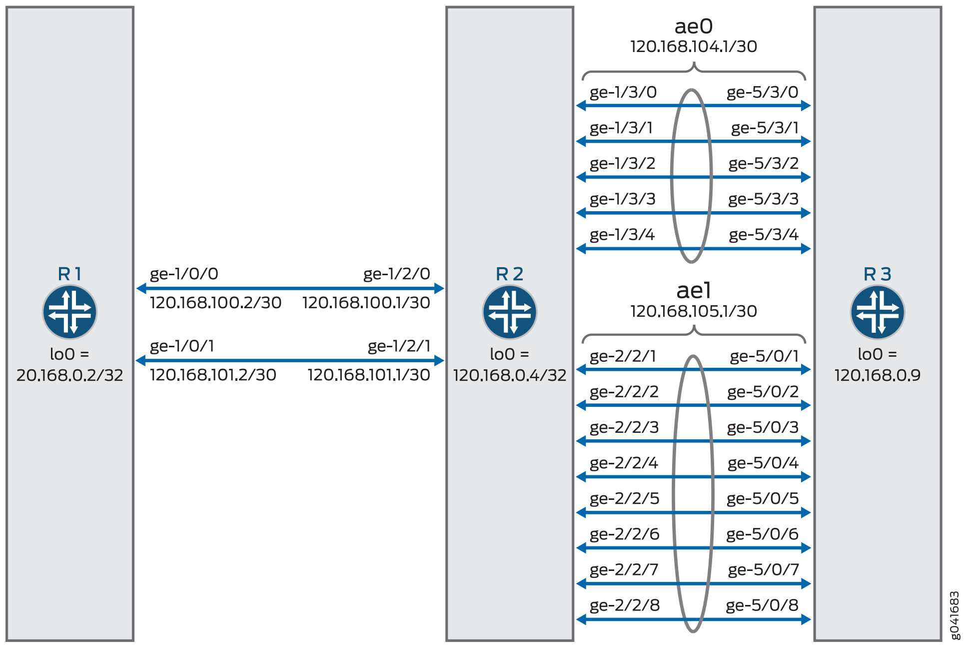

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

R1

```
set chassis aggregated-devices ethernet device-count 12
set interfaces xe-0/0/0 unit 0 family inet address 120.168.1.1/30
set interfaces xe-0/0/0 unit 0 family iso
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/1 unit 0 family inet address 120.168.2.1/30
set interfaces xe-0/0/1 unit 0 family iso
set interfaces xe-0/0/1 unit 0 family mpls
set interfaces ge-1/0/0 unit 0 family inet address 120.168.100.2/30
set interfaces ge-1/0/0 unit 0 family iso
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 family inet address 120.168.101.2/30
set interfaces ge-1/0/1 unit 0 family iso
set interfaces ge-1/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0002.00
set routing-options router-id 120.168.0.2
```

```
set routing-options autonomous-system 55
set protocols rsvp interface ge-1/0/0.0
set protocols rsvp interface ge-1/0/1.0
set protocols mpls label-switched-path videl-to-sweets to 120.168.0.9
set protocols mpls label-switched-path v-2-s-601 to 60.0.1.0
set protocols mpls label-switched-path v-2-s-601 primary v-2-s-601-primary hop-limit
  5
set protocols mpls label-switched-path v-2-s-602 to 60.0.2.0
set protocols mpls label-switched-path v-2-s-602 primary v-2-s-602-primary hop-limit
  5
set protocols mpls label-switched-path v-2-s-603 to 60.0.3.0
set protocols mpls label-switched-path v-2-s-604 to 60.0.4.0
set protocols mpls path v-2-s-601-primary 120.168.100.1 strict
set protocols mpls path v-2-s-601-primary 120.168.104.2 strict
set protocols mpls path v-2-s-602-primary 120.168.101.1 strict
set protocols mpls path v-2-s-602-primary 120.168.105.2 strict
set protocols mpls interface ge-1/0/0.0
set protocols mpls interface ge-1/0/1.0
set protocols mpls interface xe-0/0/1.0
set protocols mpls interface xe-0/0/0.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.2
set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.9
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/0/0.0
set protocols isis interface ge-1/0/1.0
set protocols isis interface lo0.0
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add
  vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-0/0/0.0
set routing-instances vpn-m5 interface xe-0/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.2:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.1.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.2.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-0/0/1.0
```

```
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-0/0/0.0
```

R2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/2/0 unit 0 family inet address 120.168.100.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 120.168.101.1/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/3/0 gigether-options 802.3ad ae0
set interfaces ge-1/3/1 gigether-options 802.3ad ae0
set interfaces ge-1/3/2 gigether-options 802.3ad ae0
set interfaces ge-1/3/3 gigether-options 802.3ad ae0
set interfaces ge-1/3/4 gigether-options 802.3ad ae0
set interfaces ge-2/2/1 gigether-options 802.3ad ae1
set interfaces ge-2/2/2 gigether-options 802.3ad ae1
set interfaces ge-2/2/3 gigether-options 802.3ad ae1
set interfaces ge-2/2/4 gigether-options 802.3ad ae1
set interfaces ge-2/2/5 gigether-options 802.3ad ae1
set interfaces ge-2/2/6 gigether-options 802.3ad ae1
set interfaces ge-2/2/7 gigether-options 802.3ad ae1
set interfaces ge-2/2/8 gigether-options 802.3ad ae1
set interfaces ae0 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.1/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.1/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
set accounting-options selective-aggregate-interface-stats disable
set protocols rsvp interface ge-1/2/0.0
set protocols rsvp interface ge-1/2/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/2/0.0
set protocols isis interface ge-1/2/1.0
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0
```

R3

```
set chassis aggregated-devices ethernet device-count 5
```

```
set interfaces xe-4/0/0 unit 0 family inet address 120.168.9.1/30
set interfaces xe-4/0/0 unit 0 family mpls
set interfaces xe-4/0/1 unit 0 family inet address 120.168.10.1/30
set interfaces xe-4/0/1 unit 0 family mpls
set interfaces ge-5/0/1 gigether-options 802.3ad ae1
set interfaces ge-5/0/2 gigether-options 802.3ad ae1
set interfaces ge-5/0/3 gigether-options 802.3ad ae1
set interfaces ge-5/0/4 gigether-options 802.3ad ae1
set interfaces ge-5/0/5 gigether-options 802.3ad ae1
set interfaces ge-5/0/6 gigether-options 802.3ad ae1
set interfaces ge-5/0/7 gigether-options 802.3ad ae1
set interfaces ge-5/0/8 gigether-options 802.3ad ae1
set interfaces ge-5/3/0 gigether-options 802.3ad ae0
set interfaces ge-5/3/1 gigether-options 802.3ad ae0
set interfaces ge-5/3/2 gigether-options 802.3ad ae0
set interfaces ge-5/3/3 gigether-options 802.3ad ae0
set interfaces ge-5/3/4 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.2/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.2/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.9/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0009.00
set routing-options router-id 120.168.0.9
set routing-options autonomous-system 55
set protocols rsvp interface xe-4/0/0.0
set protocols rsvp interface xe-4/0/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls label-switched-path to-videl to 120.168.0.2
set protocols mpls interface xe-4/0/0.0
set protocols mpls interface xe-4/0/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.9
set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.2
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add
  vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
```



```
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from protocol direct
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-4/0/0.0
set routing-instances vpn-m5 interface xe-4/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.9:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.9.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.10.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/0.0
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/1.0
```

Configuring Adaptive Load Balancing

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the R2 router:

> NOTE: Repeat this procedure for the other routers, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@R2# set aggregated-devices ethernet device-count 5
   ```

2. Configure the Gigabit Ethernet interface link connecting R2 to R1.

   ```
   [edit interfaces]
   user@R2# set ge-1/2/0 unit 0 family inet address 120.168.100.1/30
   user@R2# set ge-1/2/0 unit 0 family iso
   user@R2# set ge-1/2/0 unit 0 family mpls

   user@R2# set ge-1/2/1 unit 0 family inet address 120.168.101.1/30
   user@R2# set ge-1/2/1 unit 0 family iso
   user@R2# set ge-1/2/1 unit 0 family mpls
   ```

```
user@R2# set lo0 unit 0 family inet address 120.168.0.4/32
user@R2# set lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
```

3. Configure the five member links of the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ge-1/3/0 gigether-options 802.3ad ae0
user@R2# set ge-1/3/1 gigether-options 802.3ad ae0
user@R2# set ge-1/3/2 gigether-options 802.3ad ae0
user@R2# set ge-1/3/3 gigether-options 802.3ad ae0
user@R2# set ge-1/3/4 gigether-options 802.3ad ae0
```

4. Configure the eight member links of the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ge-2/2/1 gigether-options 802.3ad ae1
user@R2# set ge-2/2/2 gigether-options 802.3ad ae1
user@R2# set ge-2/2/3 gigether-options 802.3ad ae1
user@R2# set ge-2/2/4 gigether-options 802.3ad ae1
user@R2# set ge-2/2/5 gigether-options 802.3ad ae1
user@R2# set ge-2/2/6 gigether-options 802.3ad ae1
user@R2# set ge-2/2/7 gigether-options 802.3ad ae1
user@R2# set ge-2/2/8 gigether-options 802.3ad ae1
```

5. Enable aggregate Ethernet load balancing on ae0 of R2.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options load-balance adaptive tolerance 10
```

6. Configure the link speed for the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options link-speed 1g
```

7. Configure LACP on the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options lacp active
```

8. Configure the interface parameters for the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 unit 0 family inet address 120.168.104.1/30
user@R2# set ae0 unit 0 family iso
user@R2# set ae0 unit 0 family mpls
```

9. Enable aggregate Ethernet load balancing on ae1 of R2.

```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options load-balance adaptive tolerance 10
```

10. Configure the link speed for the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
```


user@R2# **set ae1 aggregated-ether-options link-speed 1g**

11. Configure LACP on the ae1 aggregated Ethernet bundle.

    [edit interfaces]
    user@R2# **set ae1 aggregated-ether-options lacp active**

12. Configure the interface parameters for the ae1 aggregated Ethernet bundle.

    [edit interfaces]
    user@R2# **set ae1 unit 0 family inet address 120.168.105.1/30**
    user@R2# **set ae1 unit 0 family iso**
    user@R2# **set ae1 unit 0 family mpls**

13. Disable selective aggregate Ethernet statistics.

    [edit accounting-options]
    user@R2# **set selective-aggregate-interface-stats disable**

14. Configure RSVP on all the interfaces of R2 and on the AE bundles.

    [edit protocols]
    user@R2# **set rsvp interface ge-1/2/0.0**
    user@R2# **set rsvp interface ge-1/2/1.0**
    user@R2# **set rsvp interface ae0.0**
    user@R2# **set rsvp interface ae1.0**

15. Configure MPLS on all the interfaces of R2 and on the AE bundles.

    [edit protocols]
    user@R2# **set mpls interface ge-1/2/0.0**
    user@R2# **set mpls interface ge-1/2/1.0**
    user@R2# **set mpls interface ae0.0**
    user@R2# **set mpls interface ae1.0**

16. Configure IS-IS on all the interfaces of R2 and on the AE bundles.

    [edit protocols]
    user@R2# **set isis traffic-engineering family inet shortcuts**
    user@R2# **set isis level 1 disable**
    user@R2# **set isis interface ge-1/2/0.0**
    user@R2# **set isis interface ge-1/2/1.0**
    user@R2# **set isis interface ae0.0**
    user@R2# **set isis interface ae1.0**
    user@R2# **set isis interface lo0.0**

Results

From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show accounting-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.


```
user@R2# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
  }
}
```

```
user@R2# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      address 120.168.100.1/30;
    }
    family iso;
    family mpls;
  }
}
ge-1/2/1 {
  unit 0 {
    family inet {
      address 120.168.101.1/30;
    }
    family iso;
    family mpls;
  }
}
ge-1/3/0 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/3/1 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/3/2 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/3/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/3/4 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/2/1 {
  gigether-options {
    802.3ad ae1;
  }
}
```



```
ge-2/2/2 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/3 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/4 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/6 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/7 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/8 {
  gigether-options {
    802.3ad ae1;
  }
}
ae0 {
  aggregated-ether-options {
    load-balance {
      adaptive tolerance 10;
    }
    link-speed 1g;
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 120.168.104.1/30;
    }
    family iso;
    family mpls;
  }
}
ae1 {
  aggregated-ether-options {
    load-balance {
```

```
            adaptive tolerance 10;
        }
        link-speed 1g;
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 120.168.105.1/30;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 120.168.0.4/32;
        }
        family iso {
            address 49.0001.1201.6800.0004.00;
        }
    }
}
```

```
user@R2# show accounting-options
selective-aggregate-interface-stats disable;
```

```
user@R2# show protocols
rsvp {
    interface ge-1/2/0.0;
    interface ge-1/2/1.0;
    interface ae0.0;
    interface ae1.0;
}
mpls {
    interface ge-1/2/0.0;
    interface ge-1/2/1.0;
    interface ae0.0;
    interface ae1.0;
}
isis {
    traffic-engineering {
        family inet {
            shortcuts;
        }
    }
    level 1 disable;
    interface ge-1/2/0.0;
    interface ge-1/2/1.0;
    interface ae0.0;
    interface ae1.0;
    interface lo0.0;
}
```


Verification

Confirm that the configuration is working properly.

• Verifying Adaptive Load Balancing on ae0 on page 175

Verifying Adaptive Load Balancing on ae0

Purpose Verify that packets received on the ae0 aggregated Ethernet bundle are load-balanced among the five member links.

Action From operational mode, run the **show interfaces ae0 extensive** command.

```
user@R2> show interfaces ae0 extensive
Logical interface ae0.0 (Index 325) (SNMP ifIndex 917) (Generation 134)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :        848761          9      81247024         7616
        Output:  166067308909    3503173 126900990064983  21423804256
    Adaptive Statistics:
        Adaptive Adjusts:        264
        Adaptive Scans  :      27682
        Adaptive Updates:         10
    Link:
      ge-1/3/0.0
        Input :        290888          5      29454436         3072
        Output:   33183442699     704569 25358563587277   4306031760
      ge-1/3/1.0
        Input :        162703          1      14806325          992
        Output:   33248375409     705446 25406995966732   4315342152
      ge-1/3/2.0
        Input :        127448          1      12130566          992
        Output:   33184552729     697572 25354827700261   4267192376
      ge-1/3/3.0
        Input :        121044          1      11481262         1280
        Output:   33245875402     697716 25405953405192   4265750584
      ge-1/3/4.0
        Input :        146678          1      13374435         1280
        Output:   33205071207     697870 25374651121458   4269487384
```

Meaning The member links of the ae0 aggregated Ethernet bundle are fully utilized with adaptive load balancing.

See Also • Understanding Aggregated Ethernet Load Balancing on page 158

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 128 LAG bundles on M Series, and T Series routers, and 480 LAG bundles on MX Series routers and EX9200 switches. Each LAG bundle contains up to 16 links. (Platform support depends on the Junos OS release in your installation.)

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement. For more information, see "Configuring Load Balancing on a LAG Link" on page 177.In a Layer 2 switch, one link is overutilized and other links are underutilized.

Release History Table

Release	Description
10.1	Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement.

Related Documentation

- Configuring Load Balancing on a LAG Link on page 177
- Load Balancing on a LAG Link on page 176
- *payload*

Example: Configuring Load Balancing on a LAG Link

This example configures the load-balancing hash key to use the source Layer 3 IP address option and Layer 4 header fields as well as the source and destination MAC addresses for load balancing on a link aggregation group (LAG) link:

```
[edit]
forwarding-options {
  hash-key {
    family multiservice {
      source-mac;
      destination-mac;
      payload {
        ip {
          layer-3 {
            source-ip-only;
          }
          layer-4;
        }
      }
    }
  }
}
```

NOTE: Any change in the hash key configuration requires a reboot of the FPC for the changes to take effect.

Related Documentation

- Load Balancing and Ethernet Link Aggregation on page 175
- Configuring Load Balancing on a LAG Link on page 177

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.

By default, Junos implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet.

For more information about link aggregation group (LAG) configuration, see the *Junos OS Network Interfaces Library for Routing Devices*.

Related Documentation

- Load Balancing and Ethernet Link Aggregation on page 175
- Load Balancing on a LAG Link on page 176

Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data

When multiple flows are transmitted out of an aggregated Ethernet (**ae**) interface, the flows must be distributed across the different member links evenly to enable an effective and optimal load-balancing behavior. To obtain a streamlined and robust method of load-balancing, the member link of the aggregated Ethernet interface bundle that is selected each time for load balancing plays a significant part. In Junos OS releases earlier than Release 13.2R1, on MX Series routers with Trio-based FPCs (MPCs), the selection of a member link of the **ae** interface bundle or the next-hop (or unilist of next-hops) for equal-cost multipath ECM) links is performed using a balanced mode next-hop selection methodology and an unbalanced mode of member link or next-hop selection methodology. The balanced mode of link selection uses 'n' bits in a precomputed hash value if it needs to select one of 2^n (2 raised to the power of n) next-hop in the unilist. The unbalanced mode of member-link or next-hop selection uses 8 bits in a precomputed hash to select an entry in a selector table, which is randomly done with the member link IDs of the link aggregation group (LAG) or **ae**bundle.

The term balanced versus unbalanced indicates whether a selector table is used for load balancing mechanism or not. The LAG bundle uses the unbalanced mode (selector table balancing) to balance the traffic across member links. When the traffic flows are minimal, the following problems might occur with the unbalanced mode: The link selection logic utilizes only subset bits of the precomputed hash. Regardless of the efficiency of the hashing algorithm, it is only the compressed representation of a flow. Because the inter-flow variance is very low, the resultant hashes and the subset that are computed do not provide the necessary variability to effectively utilize all the LAG member links. An excessive amount of random nature exists in the hash computation and also in the selector table. As a result, the deviation from being an optimal load-balancing technique for each child link that is selected is higher when the number of flows is lower.

The deviation per child link is defined as

Vi = ((Ci - (M/N)))/N

where

- Vi denotes the deviation for that child link 'i'.
- i denotes the child link member/index.
- Ci represents the packets transmitted for that child link 'i'.
- M signifies the total packets transmitted on that LAG bundle.
- N denotes the number of child links in that LAG.

Because of these drawbacks, for smaller number of flows, or flows with less inter-flow variance, the link utilization is skewed, and a high probability of a few child links not being utilized entirely exists. Starting with Junos OS Release 13.2R1, the capability to perform uniform load balancing and also perform rebalancing is introduced on MX Series routers with MPCs, except MPC3Es and MPC4Es. Rebalancing is not supported when load-balancing is skewed or distorted owing to a change in the number of flows.

The mechanism to record and retain states for the flows and distribute the traffic load accordingly is added. As a result, for m number of flows, they are distributed among n member links of a LAG bundle or among the unilist of next-hops in an ECMP link. This method of splitting the load among member links is called *stateful load balancing* and it uses 5-tuple information (source and destination addresses, protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

This mechanism works efficiently only for minimal number of flows (less than thousands of flows, approximately). For a larger number of flows (between 1000 and 10,000 flows), we recommend that distributed Trio-based load-balancing mechanism is used.

Consider a sample scenario in which 'n' links in the LAG are identified with link IDs of 0 through n-1. A hash table or a flow table is used to record the flows as and when they show up. The hashing key is constructed using the fields that uniquely identify a flow. The result of the lookup identifies the link_id that the flow is currently using. For each packet, the flow table based on the flow identifier is examined. If a match is found, it denotes a packet that belongs to a flow that is previously processed or detected. The link ID is associated with the flow. If a match is not found, it is the first packet that belongs to the flow. The link ID is used to select the link and the flow is inserted into the flow table.

To enable per-flow load balancing based on hash values, include the **per-flow** statement at the at the **[edit interfaces ae*X* unit *logical-unit-number* forwarding-options load-balance-stateful]** hierarchy level. By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All Packet Forwarding Engine slots are assigned the same hash value by default. To configure the load-balancing algorithm to dynamically rebalance the LAG using existing parameters, include the **rebalance *interval*** statement at the **[edit interfaces ae*X* unit *logical-unit-number* forwarding-options load-balance-stateful]** hierarchy level. This parameter periodically load balances traffic by providing a synchronized rebalance switchover across all the ingress Packet Forwarding Engines (PFEs) over a rebalance interval. You can specify the interval as a value in the range of 1 through 1000 flows per minute. To configure the load type, include the **load-type (low | medium | high)** statement at the **[edit interfaces ae*X* unit *logical-unit-number* forwarding-options load-balance-stateful]** hierarchy level.

The **stateful per-flow** option enables the load-balancing capability on AE bundles. The **rebalance** option clears the load balance state at specified intervals. The **load** option informs the Packet Forwarding Engine regarding the appropriate memory pattern to be used. If the number of flows that flow on this aggregated Ethernet interface is less (between 1 and 100 flows), then the **low** keyword can be used. Similarly for relatively higher flows (between 100 and 1000 flows), the **medium** keyword can be used and the **large** keyword can be used for the maximum flows (between 1000 and 10,000 flows). The approximate number of flows for effective load-balancing for each keyword is a derivative.

The **clear interfaces ae*X* unit *logical-unit-number* forwarding-options load-balance state** command clears the load balance state at the hardware level and enables rebalancing from the cleaned up, empty state. This clear state is triggered only when you use this command. The **clear interfaces aggregate forwarding-options load-balance state** command clears all the aggregate Ethernet interface load balancing states and re-creates them newly.

Guidelines for Configuring Stateful Load Balancing for Aggegated Ethernet Interfaces or LAG Bundles

Keep the following points in mind while configuring stateful load-balancing for aggregated Ethernet interfaces:

- When a child link is removed or added, a new aggregate selector is selected and traffic flows onto the new selector. Because the selector is empty, flows are filled in the selector. This behavior causes redistribution of flows because the old state is lost. This is the existing behavior without enabling stateful per-flow load-balancing.
- Stateful per-flow load-balancing functions on AE interfaces if the incoming traffic reaches the MPC1E, MPC2E, MPC3E-3D, MPC5E, and MPC6E line cards. Any other type of line card does not rigger this functionality. Appropriate CLI errors are displayed if the MPCs do not support this capability.

  With the ingress line card as MPC and the egress line card as MPC or DPC, this feature works properly. Stateful load-balancing is not supported if the ingress line card is a DPC and the egress line card is a DPC or an MPC.
- This capability is not supported for multicast traffic (native/flood).
- Enabling the rebalance option or clearing the load balance state can cause packet reordering for active flows because different sets of links can be selected for traffic flows.
- Although the feature performance is high, it consumes significant amount of line card memory. Approximately, 4000 logical interfaces or 16 aggregated Ethernet logical interfaces can have this feature enabled on supported MPCs. However, when the Packet Forwarding Engine hardware memory is low, depending upon the available memory, it falls back to the default load balancing mechanism. A system logging message is generated in such a situation and sent to the Routing Engine. A restriction on the number of AE interfaces that support stateful load-balancing does not exist; the limit is determined by the line cards.
- If the traffic flows become aged frequently, then the device needs to remove or refresh the load balancing states. As a result, you must configure rebalancing or run the clear command at periodic intervals for proper load-balancing. Otherwise, traffic skewing can occur. When a child link goes down or comes up, the load balancing behavior does not undergo changes on existing flows. This condition is to avoid packet reordering. New flows pick up the child link that come up. If you observe load distribution to be not very effective, you can clear the load-balancing states or use rebalancing functionality to cause an automatic clearance of the hardware states. When you configure the rebalancing facility, traffic flows can get redirected to different links, which can cause packet reordering.

Release History Table

Release	Description
13.2R1	Starting with Junos OS Release 13.2R1, the capability to perform uniform load balancing and also perform rebalancing is introduced on MX Series routers with MPCs, except MPC3Es and MPC4Es.

Related Documentation

- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces on page 181

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

The mechanism to record and retain states for the flows and distribute the traffic load accordingly is added. As a result, for m number of flows, they are distributed among n member links of a LAG bundle or among the unilist of next-hops in an ECMP link. This method of splitting the load among member links is called *stateful load balancing* and it uses 5-tuple information (source and destination addresses, protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

To configure stateful load balancing on **ae** interface bundles:

1. Specify that you want to configure an aggregated Ethernet interface.

   [edit]
   user@R2# **set interfaces ae*X* unit *logical-unit-number***

2. Specify that you want to configure stateful load-balancing.

   [edit interfaces ae*X* unit *logical-unit-number*]
   user@R2# **edit forwarding-options load-balance-stateful**

3. Enable the mechanism to perform an even, effective distribution of traffic flows across member links of an aggregated Ethernet interface (**ae**) bundle on MX Series routers with MPCs, except MPC3Es and MPC4Es.

   [edit interfaces ae*X* unit *logical-unit-number* load-balance-stateful]
   user@R2# **set per-flow**

4. Configure periodic rebalancing of traffic flows of an aggregated Ethernet bundle by clearing the load balance state at a specified interval.

   [edit interfaces ae*X* unit *logical-unit-number* load-balance-stateful]
   user@R2# **set rebalance *interval***

5. Define the load-balancing type to inform the Packet Forwarding Engine regarding the appropriate memory pattern to be used for traffic flows. The approximate number of flows for effective load-balancing for each keyword is a derivative.

   [edit interfaces ae*X* unit *logical-unit-number* load-balance-stateful]

user@R2# **set load-type (low | medium | large)**

6. Configure the address family and IP address for the **ae** interface.

   [edit interfaces]
   user@R2# **set ae0 aggregated-ether-options load-balance adaptive tolerance 10**

7. Configure the link speed for the ae0 aggregated Ethernet bundle.

   [edit interfaces ae*X* unit *logical-unit-number*]]
   user@R2# **set family *family-name* address *address***

Related Documentation

- Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data on page 178

Configuring Adaptive Load Balancing

This topic describes how to configure adaptive load balancing. Adaptive load balancing maintains efficient utilization of member link bandwidth for an aggregated Ethernet (AE) bundle. Adaptive load balancing uses a feedback mechanism to correct traffic load imbalance by adjusting the bandwidth and packet streams on links within an AE bundle.

Before you begin:

- Configure a set of interfaces with a protocol family and IP address. These interfaces can make up the membership for the AE bundle.
- Create an AE bundle by configuring a set of router interfaces as aggregated Ethernet and with a specific AE group identifier.

To configure adaptive load balancing for an AE bundles:

1. Enable adaptive load balancing on the AE bundle:

   [edit interfaces ae-*x* aggregated-ether-options load-balance]
   user@router# **set adaptive**

   NOTE: **To configure adaptive load balancing on aggregated Ethernet bundles with mixed link speeds, use the following statement:**

   user@router# **set interfaces ae0 aggregated-ether-options link-speed *mixed* load-balance *adaptive***

2. Configure the scan interval value for adaptive load balancing on the AE bundle. The scan interval value determines the length of the traffic scan by multiplying the integer value with a 30-second time period:

   [edit interfaces ae-*x* aggregated-ether-options load-balance adaptive]
   user@router# **set scan-interval *multiplier***

3. Configure the tolerance percentage value. The tolerance value determines the allowed deviation in the traffic rates among the members of the AE bundle before the router triggers an adaptive load balancing update:

   [edit interfaces ae-*x* aggregated-ether-options load-balance adaptive]
   user@router# **set tolerance** ***percentage***

4. (Optional) Enable packet-per-second-based adaptive load balancing on the AE bundle:

   [edit interfaces ae-*x* aggregated-ether-options load-balance adaptive]
   user@router# **set pps**

Related Documentation

- Understanding Aggregated Ethernet Load Balancing on page 158
- Example: Configuring Aggregated Ethernet Load Balancing on page 163
- **adaptive on page 1070**

Understanding Independent Micro BFD Sessions for LAG

Starting with Junos OS Release 13.3, this feature is supported on the following PIC/FPC types:

- PC-1XGE-XENPAK (Type 3 FPC)
- PD-4XGE-XFP (Type 4 FPC)
- PD-5-10XGE-SFPP (Type 4 FPC)
- 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 1x100GE Type 5 PICs
- All MPCs on MX Series with Ethernet MICs
- FPC-PTX-P1-A on PTX5000 with 10-Gigabit Ethernet interfaces
- FPC2-PTX-P1A on PTX5000 with 10-Gigabit Ethernet interfaces in Junos OS Release 14.1 and later
- All FPCs on PTX Series with Ethernet interfaces in Junos OS Release 14.1R3 and later 14.1 releases, and Junos 14.2 and later

TIP: See *PTX Series PIC/FPC Compatibility* for a list of PICs that are supported on each PTX Series FPC.

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. A link aggregation group (LAG) combines multiple links between devices that are in point-to-point connections, thereby increasing bandwidth, providing reliability, and allowing load balancing. To run a BFD session on LAG interfaces, configure an independent, asynchronous mode BFD session on every LAG member link in a LAG bundle. Instead of a single BFD session monitoring the status of the UDP port, independent micro BFD sessions monitor the status of individual member links.

The individual BFD sessions determine the Layer 2 and Layer 3 connectivity of each member link in the LAG. Once a BFD session is established on a particular link, the member links are attached to the LAG and the load balancer either by a static configuration or by the Link Aggregation Control Protocol (LACP). If the member links are attached to the LAG by a static configuration, the device control process acts as the client to the micro BFD session. When member links are attached to the LAG by the LACP, the LACP acts as the client to the micro BFD session.

When the micro BFD session is up, a LAG link is established and data is transmitted over that LAG link. If the micro BFD session on a member link is down, that particular member link is removed from the load balancer, and the LAG managers stop directing traffic to that link. These micro BFD sessions are independent of each other despite having a single client that manages the LAG interface.

NOTE: **Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD. Dedicated MAC mode is used by default for micro BFD sessions, in accordance with the latest draft for BFD over LAG.**

Micro BFD sessions run in the following modes:

- Distribution Mode—Micro BFD sessions are distributed by default at Layer 3.
- Non-Distribution Mode—You can configure the BFD session to run in this mode by including the **no-delegate-processing** statement under periodic packet management (PPM). In this mode, the packets are being sent or received by the Routing Engine at Layer 2.

A pair of routing devices in a LAG exchange BFD packets at a specified, regular interval. The routing device detects a neighbor failure when it stops receiving a reply after a specified interval. This allows the quick verification of member link connectivity with or without LACP. A UDP port distinguishes BFD over LAG packets from BFD over single-hop IP.

NOTE: **IANA has allocated 6784 as the UDP destination port for micro BFD.**

To enable failure detection for LAG networks for aggregated Ethernet interfaces:

- Include the **bfd-liveness-detection** statement in the configuration.
- Specify a hold-down interval value to set the minimum time that the BFD session must remain up before a state change notification is sent to the other members in the LAG network.

- Specify the minimum interval that indicates the time interval for transmitting and receiving data.
- Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to Junos OS Release 16.1, you must configure the loopback address of the remote destination as the neighbor address. Beginning with Junos OS Release 16.1, you can also configure this feature with aggregated Ethernet interface address of the remote destination as the neighbor address.

  CAUTION: Deactivate bfd-liveness-detection at the **[edit interfaces aex aggregated-ether-options]** hierarchy level or deactivate the aggregated Ethernet interface before changing the neighbor address from loopback IP address to aggregated Ethernet interface IP address. Modifying the local and neighbor address without deactivating **bfd-liveness-detection** or the aggregated Ethernet interface first might cause micro BFD sessions failure.

  NOTE: Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

NOTE: This feature works only when both the devices support BFD. If BFD is configured at one end of the LAG, this feature does not work.

For the IPv6 address family, disable duplicate address detection before configuring this feature with AE interface addresses. To disable duplicate address detection, include the **dad-disable** statement at the **[edit interface *aex* unit *y* family inet6]** hierarchy level.

Release History Table

Release	Description
16.1	Beginning with Junos OS Release 16.1, you can also configure this feature with aggregated Ethernet interface address of the remote destination as the neighbor address.
16.1	Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit.
14.1	Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to Junos OS Release 16.1, you must configure the loopback address of the remote destination as the neighbor address.
13.3	Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD.

Related Documentation

- *authentication*
- **bfd-liveness-detection on page 1095**
- *detection-time*
- *transmit-interval*
- *Configuring Independent Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG on page 186

Example: Configuring Independent Micro BFD Sessions for LAG

This example shows how to configure an independent micro BFD session for aggregated Ethernet interfaces.

Requirements

This example uses the following hardware and software components:

- MX Series routers with Junos Trio chipset
- T Series routers with Type 4 FPC or Type 5 FPC

  BFD for LAG is supported on the following PIC types on T-Series:

  - PC-1XGE-XENPAK (Type 3 FPC),
  - PD-4XGE-XFP (Type 4 FPC),

    - PD-5-10XGE-SFPP (Type 4 FPC),
    - 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 1X100GE Type 5 PICs
- PTX Series routers with 24X10GE (LAN/WAN) SFPP
- Junos OS Release 13.3 or later running on all devices

Overview

The example includes two routers that are directly connected. Configure two aggregated Ethernet interfaces, AE0 for IPv4 connectivity and AE1 for IPv6 connectivity. Configure micro BFD session on the AE0 bundle using IPv4 addresses as local and neighbor endpoints on both routers. Configure micro BFD session on the AE1 bundle using IPv6 addresses as local and neighbor endpoints on both routers. This example verifies that independent micro BFD sessions are active in the output.

Topology

Figure 10 on page 187 shows the sample topology.

Figure 10: Configuring an Independent Micro BFD Session for LAG

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R0

```
set interfaces ge-1/0/1 unit 0 family inet address 20.20.20.1/30
set interfaces ge-1/0/1 unit 0 family inet6 address 3ffe::1:1/126
set interfaces xe-4/0/0 gigether-options 802.3ad ae0
set interfaces xe-4/0/1 gigether-options 802.3ad ae0
set interfaces xe-4/1/0 gigether-options 802.3ad ae1
set interfaces xe-4/1/1 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.107/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval
  100
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.102
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.107
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 10.0.0.1/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval
  100
```

```
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet6 address 5555::1/126
set interface ae1 unit 0 family inet6 dad-disable
set routing-options nonstop-routing
set routing-options static route 30.30.30.0/30 next-hop 10.0.0.2
set routing-options rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
set protocols bfd traceoptions file bfd
set protocols bfd traceoptions file size 100m
set protocols bfd traceoptions file files 10
set protocols bfd traceoptions flag all
```

Router R1

```
set interfaces ge-1/1/8 unit 0 family inet address 30.30.30.1/30
set interfaces ge-1/1/8 unit 0 family inet6 address 3ffe::1:2/126
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces xe-0/0/2 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.102/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::bb:bb:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval
  150
set interfaces ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.107
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.102
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family inet address 10.0.0.2/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval
  200
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp passive
set interfaces ae1 unit 0 family inet6 address 5555::2/126
set routing-options static route 20.20.20.0/30 next-hop 10.0.0.1
set routing-options rib inet6.0 static route 3ffe::1:1/126 next-hop 5555::1
```


Configuring a Micro BFD Session for Aggregated Ethernet Interfaces

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see "*Using the CLI Editor in Configuration Mode*" in the *CLI User Guide*.

NOTE: Repeat this procedure for Router R1, modifying the appropriate interface names, addresses, and any other parameters for each router.

To configure a micro BFD session for aggregated Ethernet interfaces on Router R0:

1. Configure the physical interfaces.

   ```
   [edit interfaces]
   user@R0# set ge-1/0/1 unit 0 family inet address 20.20.20.1/30
   user@R0# set ge-1/0/1 unit 0 family inet6 address 3ffe::1:1/126
   user@R0# set xe-4/0/0 gigether-options 802.3ad ae0
   user@R0# set xe-4/0/1 gigether-options 802.3ad ae0
   user@R0# set xe-4/1/0 gigether-options 802.3ad ae1
   user@R0# set xe-4/1/1 gigether-options 802.3ad ae1
   ```

2. Configure the loopback interface.

   ```
   [edit interfaces]
   user@R0# set lo0 unit 0 family inet address 10.255.106.107/32
   user@R0# set lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/128
   ```

3. Configure an IP address on the aggregated Ethernet interface ae0 with either IPv4 or IPv6 addresses, as per your network requirements.

   ```
   [edit interfaces]
   user@R0# set ae0 unit 0 family inet address 10.0.0.1/30
   ```

4. Set the routing option, create a static route, and set the next-hop address.

   NOTE: You can configure either an IPv4 or IPv6 static route, depending on your network requirements.

   ```
   [edit routing-options]
   user@R0# set nonstop-routing
   user@R0# set static route 30.30.30.0/30 next-hop 10.0.0.2
   user@R0# set rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
   ```

5. Configure the Link Aggregation Control Protocol (LACP).

   ```
   [edit interfaces]
   user@R0# set ae0 aggregated-ether-options lacp active
   ```

6. Configure BFD for the aggregated Ethernet interface ae0, and specify the minimum interval, local IP address, and the neighbor IP address.

```
[edit interfaces]
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection
  minimum-interval 100
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection neighbor
  10.255.106.102
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.107
user@R0# set ae0 aggregated-ether-options minimum-links 1
user@R0# set ae0 aggregated-ether-options link-speed 10g
```

7. Configure an IP addresse on the aggregated Ethernet interface ae1.

   You can assign either IPv4 or IPv6 addresses as per your network requirements.

```
[edit interfaces]
user@R0# set ae1 unit 0 family inet6 address 5555::1/126
```

8. Configure BFD for the aggregated Ethernet interface ae1.

```
[edit interfaces]
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection
  minimum-interval 100
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection neighbor
  201:DB8:251::bb:bb:1
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection local-address
  201:DB8:251::aa:aa:1
user@R0# set ae1 aggregated-ether-options minimum-links 1
user@R0# set ae1 aggregated-ether-options link-speed 10g
```

NOTE: Beginning with Junos OS Release 16.1, you can also configure this feature with the AE interface address as the local address in a micro BFD session.

Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

9. Configure tracing options for BFD for troubleshooting.

```
[edit protocols]
user@R0# set bfd traceoptions file bfd
user@R0# set bfd traceoptions file size 100m
user@R0# set bfd traceoptions file files 10
user@R0# set bfd traceoptions flag all
```

Results

From configuration mode, enter the **show interfaces**, **show protocols**, and **show routing-options** commands and confirm your configuration. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0> show interfaces
traceoptions {
  flag bfd-events;
}
ge-1/0/1 {
  unit 0 {
    family inet {
      address 20.20.20.1/30;
    }
    family inet6 {
      address 3ffe::1:1/126;
    }
  }
}
xe-4/0/0 {
  enable;
  gigether-options {
    802.3ad ae0;
  }
}
xe-4/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
xe-4/1/0 {
  enable;
  gigether-options {
    802.3ad ae1;
  }
}
xe-4/1/1 {
  gigether-options {
    802.3ad ae1;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.106.107/32;
    }
    family inet6 {
      address 201:DB8:251::aa:aa:1/128;
    }
  }
}
ae0 {
  aggregated-ether-options {
```

```
        bfd-liveness-detection {
            minimum-interval 100;
            neighbor 10.255.106.102;
            local-address 10.255.106.107;
        }
        minimum-links 1;
        link-speed 10g;
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 10.0.0.1/30;
        }
    }
}
ae1 {
    aggregated-ether-options {
        bfd-liveness-detection {
            minimum-interval 100;
            multiplier 3;
            neighbor 201:DB8:251::bb:bb:1;
            local-address 201:DB8:251::aa:aa:1;
        }
        minimum-links 1
        link-speed 10g;
    }
    unit 0 {
        family inet6 {
            address 5555::1/126;
        }
    }
}
```

```
user@R0> show protocols
bfd {
    traceoptions {
        file bfd size 100m files 10;
        flag all;
    }
}
```

```
user@R0> show routing-options
nonstop-routing ;
    rib inet6.0 {
    static {
        route 3ffe:1:2/126 {
            next-hop 5555::2;
        }
    }
}
static {
    route 30.30.30.0/30 {
        next-hop 10.0.0.2;
    }
}
```


If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

Verification

Confirm that the configuration is working properly.

- Verifying That the Independent BFD Sessions Are Up on page 193
- Viewing Detailed BFD Events on page 195

Verifying That the Independent BFD Sessions Are Up

Purpose Verify that the micro BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R0> show bfd session extensive
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.255.106.102                Up        xe-4/0/0       9.000     3.000        3

 Client LACPD, TX interval 0.100, RX interval 0.100
 Session up time 4d 23:13, previous down time 00:00:06
 Local diagnostic None, remote diagnostic None
 Remote heard, hears us, version 1
 Replicated
 Session type: Micro BFD
 Min async interval 0.100, min slow interval 1.000
 Adaptive async TX interval 0.100, RX interval 0.100
 Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
 Local discriminator 21, remote discriminator 75
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x0

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.255.106.102                Up        xe-4/0/1       9.000     3.000        3

 Client LACPD, TX interval 0.100, RX interval 0.100
 Session up time 4d 23:13, previous down time 00:00:07
 Local diagnostic None, remote diagnostic None
 Remote heard, hears us, version 1
 Replicated
 Session type: Micro BFD
 Min async interval 0.100, min slow interval 1.000
 Adaptive async TX interval 0.100, RX interval 0.100
 Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
 Local discriminator 19, remote discriminator 74
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x0

                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
201:DB8:251::bb:bb:1            Up       xe-4/1/1       9.000     3.000
      3
 Client LACPD, TX interval 0.100, RX interval 0.100
 Session up time 4d 23:13
 Local diagnostic None, remote diagnostic None
 Remote not heard, hears us, version 1
 Replicated
 Session type: Micro BFD
 Min async interval 0.100, min slow interval 1.000
 Adaptive async TX interval 0.100, RX interval 0.100
 Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
 Local discriminator 17, remote discriminator 67
 Echo mode disabled/inactive, no-absorb, no-refresh
 Remote is control-plane independent
  Session ID: 0x0
```



```
                                                  Detect    Transmit
Address                  State     Interface      Time      Interval  Multiplier
201:DB8:251::bb:bb:1               UP        xe-4/1/0       9.000     3.000
      3
 Client LACPD, TX interval 0.100, RX interval 0.100
 Session up time 4d 23:13
 Local diagnostic None, remote diagnostic None
 Remote not heard, hears us, version 1
 Replicated
 Session type: Micro BFD
 Min async interval 0.100, min slow interval 1.000
 Adaptive async TX interval 0.100, RX interval 0.100
 Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
 Local discriminator 16, remote discriminator 66
 Echo mode disabled/inactive, no-absorb, no-refresh
 Remote is control-plane independent
  Session ID: 0x0

4 sessions, 4 clients
Cumulative transmit rate 2.0 pps, cumulative receive rate 1.7 pps
```

Meaning The Micro BFD field represents the independent micro BFD sessions running on the links in a LAG. The TX interval *item*, RX interval *item* output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if required.

Action From operational mode, enter the **file show /var/log/bfd** command.

```
user@R0> file show /var/log/bfd
Jun  5 00:48:59    Protocol (1) len 1: BFD
Jun  5 00:48:59    Data (9) len 41: (hex) 42 46 44 20 6e 65 69 67 68 62 6f 72 20
 31 30 2e 30 2e 30
Jun  5 00:48:59 PPM Trace: BFD neighbor 10.255.106.102 (IFL 349) set, 9 0
Jun  5 00:48:59 Received Downstream RcvPkt (19) len 108:
Jun  5 00:48:59    IfIndex (3) len 4: 329
Jun  5 00:48:59    Protocol (1) len 1: BFD
Jun  5 00:48:59    SrcAddr (5) len 8: 10.255.106.102
Jun  5 00:48:59    Data (9) len 24: (hex) 00 88 03 18 00 00 00 4b 00 00 00 15 00
 2d c6 c0 00 2d c6
Jun  5 00:48:59    PktError (26) len 4: 0
Jun  5 00:48:59    RtblIdx (24) len 4: 0
Jun  5 00:48:59    MultiHop (64) len 1: (hex) 00
Jun  5 00:48:59    Unknown (168) len 1: (hex) 01
Jun  5 00:48:59    Unknown (171) len 2: (hex) 02 3d
Jun  5 00:48:59    Unknown (172) len 6: (hex) 80 71 1f c7 81 c0
Jun  5 00:48:59    Authenticated (121) len 1: (hex) 01
Jun  5 00:48:59 BFD packet from 10.0.0.2 (IFL 329), len 24
Jun  5 00:48:59    Ver 0, diag 0, mult 3, len 24
Jun  5 00:48:59    Flags: IHU Fate
```

```
Jun  5 00:48:59    My discr 0x0000004b, your discr 0x00000015
Jun  5 00:48:59    Tx ivl 3000000, rx ivl 3000000, echo rx ivl 0
Jun  5 00:48:59 [THROTTLE]bfdd_rate_limit_can_accept_pkt: session 10.255.106.102
 is up or already in program thread
Jun  5 00:48:59 Replicate: marked session (discr 21) for update
```

Meaning BFD messages are being written to the specified trace file.

Related Documentation

- *authentication*
- bfd-liveness-detection on page 1095
- *detection-time*
- *Configuring Independent Micro BFD Sessions for LAG*
- Understanding Independent Micro BFD Sessions for LAG on page 183

Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces

T Series and TX Matrix routers support multicast statistics collection on aggregated Ethernet interfaces in both ingress and egress directions. The multicast statistics functionality can be configured on a physical interface thus enabling multicast accounting for all the logical interfaces below the physical interface.

The multicast statistics information is displayed only when the interface is configured with the **multicast-statistics** statement, which is not enabled by default.

Multicast statistics collection requires at least one logical interface is configured with family inet or inet6; otherwise, the commit for **multicast-statistics** will fail.

The multicast in/out statistics can be obtained via interfaces statistics query through CLI and via MIB objects through SNMP query.

To configure multicast statistics:

1. Include the **multicast-statistics** statement at the **[edit interfaces interface-name]** hierarchy level.

An example of a multicast statistics configuration for an aggregated Ethernet interface follows:

```
[edit interfaces]
ae0 {
  multicast-statistics;
}
```

To display multicast statistics, use the **show interfaces** ***interface-name*** **statistics detail** command.

Related Documentation

- *multicast-statistics*
- Configuring Multicast Statistics Collection on Ethernet Interfaces on page 22
- *Ethernet Interfaces Feature Guide for Routing Devices*

Deleting an Aggregated Ethernet Interface

There are two approaches to deleting an aggregated Ethernet interface:

- You can delete an aggregated Ethernet interface from the interface configuration. The Junos OS removes the configuration statements related to **ae***x* and sets this interface to down state.
- You can also permanently remove the aggregated Ethernet interface from the device configuration by deleting it from the device-count on the routing device.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

   This step changes the interface state to down and removing the configuration statements related to **ae***x*.

   ```
   [edit]
   user@host# delete interfaces aex
   ```

2. Delete the interface from the device count.

   ```
   [edit]
   user@host# delete chassis aggregated-devices ethernet device-count
   ```

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 108
- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 133
- Aggregated Ethernet Interfaces Overview on page 102
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for Link Aggregation Control Protocol (LACP) packets only.

BEST PRACTICE: **We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.**

This topic describes:

- Disabling or Enabling Distributed Periodic Packet Management Globally on page 198
- Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 198

Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
user@switch# delete ppm no-delegate-processing
```

Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

Related Documentation

- *Understanding Distributed Periodic Packet Management*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*

ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces Overview

Starting with Junos OS Relase 16.1R1, you can configure ITU-T Y.1731 standard-compliant Ethernet loss measurement (ETH-LM), Ethernet synthetic loss measurement (ETH-SLM), and Ethernet delay measurement (ETH- DM) capabilities on aggregated Ethernet (ae) interfaces. These ITU-T Y.1731 OAM services or performance monitoring techniques can be measured by on-demand mode (triggered through the CLI) or by proactive mode (triggered by the iterator application). These performance monitoring functionalities are supported on the following platforms:

- MX Series routers with 16-port 10-Gigabit Ethernet MPCs and Trio-based FPCs (MPCs), where the same level of support for the Ethernet services OAM mechanisms on non-aggregated Ethernet interfaces is available on AE interfaces
- MX2020 routers
- ETH-DM is supported on MPC3E and MPC4E modules with only software timestamping
- ETH-SLM is supported on MPC3E and MPC4E modules.

Also, connectivity fault management (CFM) sessions established on the AE interfaces can be distributed to the Packet Forwarding Engine, apart from being handled on the Routing engine. This capability to distribute CFM sessions is useful in both scaled topologies and gracful Routing Engine switchover (GRES) for CFM sessions.

Connectivity fault management (CFM) sessions operate in centralized mode over AE interfaces by default. Y.1731 performance monitoring (PM) is supported on centralized CFM sessions over AE interfaces. Also, distribution of CFM session over AE interfaces to line cards is supported from Junos OS Release 13.3. To enable the distribution of CFM sessions and to operate in centralized mode, include the **ppm delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level. The mechanism that enables distribution of CFM sessions over AE interfaces provides the underlying infrastructure to support PM over AE interfaces. In addition, periodic packet management (PPM) handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With PPM processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run performance monitoring processes on the Packet Forwarding Engine.

For Ethernet delay measurement, hardware-assisted timestamping is supported on AE interfaces, similar to the support that exists on non-AE interfaces. Only hardware-based timestamping is supported because it is performed in the received path of the protocol data unit (PDU) packets, whereas software-based timestamping needs to be performed on the transmitted path and is not supported. For software timestamping, ETH-DM PDUs need to be transmitted and received on the same line card (same member of the AE interface). All the received ETH-DM PDUs are always redirected to the anchor Packet Forwarding Engine. In the transmission path, if the interface on the anchor Packet Forwarding Engine goes down, then the OAM pdus are redirected to one of the subordinate or member FPCs. Therefore, the processing of ETH-DM PDUs always occurs at the CPU of the line card or module that hosts the anchor Packet Forwarding Engine. ETH-DM is supported on AE interfaces with CCC, bridge, virtual private LAN service (VPLS), and inet address families. ETH-DM is supported for both active-active and active-standby modes of AE interfaces. For one-way delay measurement (1DM), the system clocks of the initiator MEP that transmits a request frame and the responder MEP that receives a reply frame need to be synchronized.

For Ethernet loss measurement on AE interfaces, with the active-standby mode of the interfaces, transmission and reception of PDUs is always through the Packet Forwarding Engine that hosts the active link. For the active-standby mode of the AE interfaces, you can configure a maximum of only two member links. ETH-LM is supported only when all the active member or child links are on the same Packet Forwarding Engine. For the downstream maintenance endpoints (MEPs), ETH-LM is supported for CCC, VPLS, and bridge address families, and for upward MEPs, ETH-LM is supported only for CCC families. In the transmission path, with active-standby links of AE interfaces, whenever the active child link fails, if the standby link is non-local, the packets are redirected to the new active link. When this redirection occurs, the ETH-LM counters are reset. If the standby link is on same Packet Forwarding Engine as the active link, then the counters are not reset because the counters are read on the local Packet Forwarding Engine memory and to prevent the other end of the session to treat new Packet Forwarding Engine counters as losses owing to reset of the counters. In the received path, with active-standby links of AE interfaces, all the child links are programmed in the input list using next-hops to redirect the packets to the anchor FPC after copying the counters in the Packet Forwarding Engine. For Ethernet synthetic loss measurement (SLM), processing of SLM PDUs for requests and responses similar to other protocols from the line card CPU is implemented. All other computation and data are software-based. ETH-SLM is supported on AE interfaces for CCC, bridge, VPLS, and inet families.

NOTE: Starting with Junos OS Release 16.1, Ethernet loss measurement over an aggregated Ethernet (ae) interface is not supported when the enhanced LAG functionality is enabled on a router. The enhanced LAG capability is enabled by default when you configure enhanced IP services mode by including the **network-services enhanced-ip** statement at the **[edit chassis]** hierarchy level. For Ethernet loss measurement to work properly, you must disable the enhanced LAG functionality by entering the **set chassis aggregated-devices disable-lag-enhanced** statement. Also, CFM is not supported when enhanced LAG is enabled.

Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the **performance-monitoring** statement and its substatements at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level) is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

Before you can start an ETH-DM, ETH-LM, or ETH-SLM measurement sessions across an aggregated Ethernet service, you must configure two MX Series routers to support these measurement sessions. On each router, configure two physical or logical AE interfaces connected by a VLAN by including the **interface ae-*fpc/pic/port* unit *logical-unit-number*** **vlan-id *vlan-id*** statement at the **[edit interfaces]** hierarchy level and on each router, attach the peer MEPs to the interfaces by including the **mep mep-id interface *interface-name* (protect | working)** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]** hierarchy level.

Release History Table

Release	Description
16.1R1	Starting with Junos OS Relase 16.1R1, you can configure ITU-T Y.1731 standard-compliant Ethernet loss measurement (ETH-LM), Ethernet synthetic loss measurement (ETH-SLM), and Ethernet delay measurement (ETH- DM) capabilities on aggregated Ethernet (ae) interfaces.
16.1	Starting with Junos OS Release 16.1, Ethernet loss measurement over an aggregated Ethernet (ae) interface is not supported when the enhanced LAG functionality is enabled on a router.
16.1	Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the **performance-monitoring** statement and its substatements at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level) is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

Related Documentation

-

Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet Interfaces

Keep the following points in mind while you configure ETH-LM, ETH-SLM, and ETH-DM capabilities on aggregated Ethernet (ae-) interfaces:

- The scaling limits and performance considerations for distributed periodic packet management (PPM) sessions. The scaling limits for distributed PPM sessions over aggregated Ethernet (AE) interfaces are identical to the maximum supported numbers for continuity check messages (CCM) over AE interfaces.
- SLA iterators always coexist with CCM sessions. Therefore, while configuring a scaled environment, you must account for CCM sessions should be accounted along with SLA iterators. The following table describes the maximum number of distributed PM sessions you can configure for different CCM intervals per line card and per router (system-wide value).
- A mixed operation of distributed and centralized modes for performance monitoring (PM) sessions is not supported on AE interfaces, if the interfaces that form the aggregated Ethernet bundle are in mixed mode.
- The limitations for performance monitoring (PM) capabilities for non-AE interfaces apply equally well for AE interfaces. For example, flapping of sessions resets the PM statistics.

- The limitations that exist with distributed PPM sessions are valid for performance monitoring capabilities over AE interfaces because measurements are always performed on CCM sessions.
- For ETH-LM over AE interfaces in an active-standby setup, if active and standby line cards are swapped, then the measurements during this window are ignored.
- For ETH-DM over AE interfaces, the additional time that is taken for packet transmission (packets are redirected to anchor in the received [Rx] direction and to the active child FPC in the transmitted [Tx] direction) is computed in the delay measurement.
- For ETH-LM over AE interfaces, in an active-standby setup, whenever the link failover from the active interface to the standby interface happens, the counters are reset.

Related Documentation

-

Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links

This document provides an overview of targeted distribution of static logical interfaces across aggregated Ethernet member links, and an example for configuring targeted distribution.

- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links Overview on page 202
- Example: Configuring Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links on page 203

Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links Overview

Targeted distribution provides a mechanism to direct traffic through specified links for an aggregated Ethernet bundle, and also assigns roles to member links to handle link failure scenarios.

The targeted distribution of static logical interfaces is also used to accurately enforce egress class-of-service (CoS) profiles. Without this feature, the enforcement of egress class-of-service profiles is distributed among the individual member interface schedulers, shapers, or policers instantiated in each Packet Forwarding Engine that hosts a member link. In the absence of targeted distribution of aggregated Ethernet bundles, traffic destined through a logical interface of a bundle can exit through any of the members based on the hashing algorithm. As a result, it is not possible to determine which link is used to forward traffic. Distributed egress policy enforcement relies on traffic load balancing, which might not be accurate all the time. With targeted distribution, logical interface traffic is directed to a certain member link or a number of member links. Targeted distribution ensures an accurate policy enforcement that is no longer distributed for a given logical interface.

NOTE: Irrespective of the family configured for the logical interface, the targeted distribution feature is applicable to both Layer 3 and Layer 2 interfaces.

You can form distribution lists consisting of member links of the aggregated Ethernet bundle and you can assign roles to these lists, such as primary, backup, and standby. A distribution list specified as primary ensures that traffic is load-balanced among all the links in the primary list. If all links within the primary list are up, traffic is forwarded on those links. If some of the links within a primary list fail, the remaining links carry traffic. If all links within the primary list go down, only then the links in the backup list start carrying traffic. If some of links within the backup list fail, the remaining links in the backup list carry traffic. If all the links within the primary list and the backup list go down, only then the links in the standby list start carrying traffic. When the primary member links come back online, they continue to carry traffic.

You can configure distribution lists for primary links and backup links. The remaining links are added to a defined standby list. You can make changes to the distribution lists and their roles by configuring them again. When targeted distribution lists are deleted, all links carry traffic. When you commit the configuration, the member links are assigned the specified roles irrespective of whether the links are up or down.

NOTE: The feature is supported only on MX Series routers with MPCs, and with the **enhanced-ip** configuration enabled.

The outbound traffic of a Layer 3 host is distributed among all the member links of an aggregated Ethernet bundle. Targeted distribution is implemented only for the transit traffic.

Example: Configuring Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links

This example shows how to configure targeted distribution lists for aggregated Ethernet member links as primary, or backup. Member links are assigned membership to the distribution lists. Logical interfaces of the aggregated Ethernet are then assigned membership to the primary list and the backup list.

- Requirements on page 203
- Overview on page 204
- Configuration on page 204
- Verification on page 207

Requirements

This example uses the following software and hardware components:

- Junos OS Release 16.1 and later releases

- One MX Series 3D Universal Edge Router

Overview

Targeted distribution provides a mechanism to direct traffic through specified links of an aggregated Ethernet bundle, and also assigns roles to member links to handle link failure scenarios.

You can configure targeted distribution to load-balance the traffic between the aggregated Ethernet bundle member links. You can map a logical interface to a single link only for the outgoing traffic.

This example uses the **apply-groups** configuration for specifying the distribution lists for the logical interfaces of the aggregated Ethernet member links. You can use the apply-groups statement to inherit the Junos OS configuration statements from a configuration group. The apply-groups configuration statement in the example shows the odd numbered member links of the aggregated Ethernet bundle being assigned the primary list dl2 and even numbered member links being assigned primary list dl1.

The aggregated Ethernet interface used in this example is ae10 with units 101, 102, 103, and 104. The Gigabit Ethernet interface, ge-0/0/3 is specified as distribution list dl1 and ge-0/0/4 as dl2. The logical interface unit numbers of the aggregated Ethernet bundle ending in an odd number are assigned to the distribution list *dl1* as the primary list, and those ending in an even number are assigned the distribution list *dl2* as the primary list.

To configure targeted distribution you must:

1. Create a global apply group.
2. Assign each member of the aggregated Ethernet interface to a different distribution list.
3. Attach the apply group to the aggregated Ethernet interface.
4. Create the logical interfaces. The apply group automatically assigns the distribution lists to each member of the aggregated Ethernet bundle as required.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
[edit groups GR-AE-ACCESS-DISTRIBUTION]
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description “matched-odd”
  targeted-distribution primary-list dl2
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description “matched-odd”
  targeted-distribution backup-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description “matched-even”
  targeted-distribution primary-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description “matched-even”
  targeted-distribution backup-list dl2
```



```
user@host# set interfaces ge-0/0/3 apply-groups-except INTF gigether-options 802.3ad
  ae10 distribution-list dl1
user@host# set interfaces ge-0/0/4 apply-groups-except INTF gigether-options 802.3ad
  ae10 distribution-list dl2
user@host# set interfaces <ae*> apply-groups GR-AE-ACCESS-DISTRIBUTION
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 101 vlan-id 101 family inet address 10.1.0.1/16
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 102 vlan-id 102 family inet address 10.2.0.1/16
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 103 vlan-id 103 family inet address 10.3.0.1/16
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 104 vlan-id 104 family inet address 10.4.0.1/16
```

To configure targeted distribution:

1. Create a global apply group and specify the primary list and the backup list.

```
[edit groups GR-AE-ACCESS-DISTRIBUTION]
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description "matched-odd"
  targeted-distribution primary-list dl2
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description "matched-odd"
  targeted-distribution backup-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description "matched-even"
  targeted-distribution primary-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description "matched-even"
  targeted-distribution backup-list dl2
```

2. Assign each member of the aggregated Ethernet bundle to a different distribution list.

```
[edit]
user@host# set interfaces ge-0/0/3 apply-groups-except INTF gigether-options
  802.3ad ae10 distribution-list dl1
[edit]
user@host# set interfaces ge-0/0/4 apply-groups-except INTF gigether-options
  802.3ad ae10 distribution-list dl2
```

3. Attach the defined apply group to the aggregated Ethernet interface.

```
[edit]
user@host# set interfaces ae10 apply-groups GR-AE-ACCESS-DISTRIBUTION
```

4. Create the logical interfaces and configure its provisions.

```
[edit]
user@host# set interfaces ae10 apply-groups GR-AE-ACCESS-DISTRIBUTION
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation
  flexible-ethernet-services set unit 101 vlan-id 101 family inet address 10.1.0.1/16
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 102 vlan-id 102 family inet address 10.2.0.1/16
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 103 vlan-id 103 family inet address 10.3.0.1/16
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation
  flexible-ethernet-services unit 104 vlan-id 104 family inet address 10.4.0.1/16
```

Results From configuration mode, confirm your configuration by using the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show groups GR-AE-ACCESS-DISTRIBUTION
interfaces {
  <ae*> {
    unit "<*[1 3 5 7 9]>" {
      description "matched odd";
      targeted-distribution {
        primary-list dl2;
        backup-list dl1;
      }
    }
    unit "<*[0 2 4 6 8]>" {
      description "matched even";
      targeted-distribution {
        primary-list dl1;
        backup-list dl2;
      }
    }
  }
}
```

```
user@host# show interfaces ge-0/0/3
apply-groups-except INTF;
gigether-options {
  802.3ad {
    ae10;
    distribution-list dl1;
  }
}
```

```
user@host# show interfaces ge-0/0/4
apply-groups-except INTF;
gigether-options {
  802.3ad {
    ae10;
    distribution-list dl2;
  }
}
```

```
user@host# show interfaces ae10 apply-groups
apply-groups GR-AE-ACCESS-DISTRIBUTION;
```

```
user@host# show interfaces ae10
apply-groups GR-AE-ACCESS-DISTRIBUTION;
flexible-vlan-tagging; encapsulation flexible-ethernet-services;
unit 101 {
  vlan-id 101;
  family inet {
    address 10.1.0.1/16 {
    }
  }
}
unit 102 {
  vlan-id 102;
```



```
    family inet {
      address 10.2.0.1/16 {
      }
    }
  }
  unit 103 {
    vlan-id 103;
    family inet {
      address 10.3.0.1/16 {
      }
    }
  }
  unit 104 {
    vlan-id 104;
    family inet {
      address 10.4.0.1/16 {
      }
    }
  }
```

Verification

Verifying Targeted Distribution of Logical Interfaces

Purpose Verify that the logical interfaces are assigned to the distribution lists.

Action To verify that the logical interfaces are assigned to the distribution lists, enter the **show interfaces detail or extensive** command.

The **show interfaces detail or extensive** command output shows the logical interfaces ending in an odd number being assigned to the distribution list, *dl1* (**ge-0/0/3**) and those ending in an even number being assigned distribution list, *dl2* (**ge-0/0/4**) by default. If there is a failure of either of those interfaces, the logical interfaces switch to the interfaces in the backup list or continue to use the active member interface. For example, on the aggregated Ethernet bundle, **ae10.101**, the primary interface shown is **ge-0/0/4** and on the aggregated Ethernet bundle **ae10.102**, the primary interface is **ge-0/0/3** and similarly for the other logical interfaces.

```
user@host# run show interfaces extensive ae10

Physical interface: ae10, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 612, Generation: 132
  Link-level type: Flexible-Ethernet, MTU: 9000, Speed: 2Gbps, BPDU Error: None,
 MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:05:86:1e:70:c1, Hardware address: 00:05:86:1e:70:c1
  Last flapped   : 2016-08-30 16:15:28 PDT (00:43:15 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                77194                  200 bps
```

```
  Input  packets:                   0                    0 pps
  Output packets:                 300                    0 pps
  IPv6 transit statistics:
   Input  bytes  :                  0
   Output bytes  :                  0
   Input  packets:                  0
   Output packets:                  0
 Dropped traffic statistics due to STP State:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
 Input errors:
   Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
 Output errors:
   Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
 Ingress queues: 8 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0                             0                    0                    0

   1                             0                    0                    0

   2                             0                    0                    0

   3                             0                    0                    0

 Egress queues: 8 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0                             0                    0                    0

   1                             0                    0                    0

   2                             0                    0                    0

   3                             0                    0                    0

 Queue number:         Mapped forwarding classes
   0                   best-effort
   1                   expedited-forwarding
   2                   assured-forwarding
   3                   network-control

 Logical interface ae10.101 (Index 345) (SNMP ifIndex 617) (Generation 154)
   Description: matched odd
   Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.101 ]  Encapsulation: ENET2
   Statistics        Packets        pps         Bytes          bps
   Bundle:
       Input :             0          0             0            0
       Output:             2          0            92            0
   Adaptive Statistics:
       Adaptive Adjusts:         0
       Adaptive Scans  :         0
       Adaptive Updates:         0
   Link:
     ge-0/0/3.101
       Input :             0          0             0            0
       Output:             2          0            92            0
```



```
      ge-0/0/4.101
        Input :            0          0             0             0
        Output:            0          0             0             0


    Aggregate member links: 2

    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-0/0/3.101               0           0            0            0
      ge-0/0/4.101               0           0            0            0

    List-Type       Status
    Primary         Active
                Interfaces:
                ge-0/0/4        Up
    List-Type       Status
    Backup          Waiting
                Interfaces:
                ge-0/0/3        Up
    List-Type       Status
    Standby         Down

    Protocol inet, MTU: 8978, Generation: 198, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.1.0.1/15, Local: 10.1.0.2, Broadcast: 10.1.0.3, Generation:
154
    Protocol multiservice, MTU: Unlimited, Generation: 199, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae10.102 (Index 344) (SNMP ifIndex 615) (Generation 153)
    Description: matched even
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.102 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :            0          0             0             0
        Output:            4          0           296             0
    Adaptive Statistics:
        Adaptive Adjusts:         0
        Adaptive Scans  :         0
        Adaptive Updates:         0
    Link:
      ge-0/0/3.102
        Input :            0          0             0             0
        Output:            4          0           296             0
      ge-0/0/4.102
        Input :            0          0             0             0
        Output:            0          0             0             0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-0/0/3.102               0           0            0            0
      ge-0/0/4.102               0           0            0            0

    List-Type       Status
    Primary         Active
                Interfaces:
                ge-0/0/3        Up
    List-Type       Status
    Backup          Waiting
                Interfaces:
                ge-0/0/4        Up
    List-Type       Status
```

```
    Standby         Down

    Protocol inet, MTU: 8978, Generation: 196, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.2.0.1 , Local: 10.2.0.1, Broadcast: 10.2.0.3, Generation:
152
    Protocol multiservice, MTU: Unlimited, Generation: 197, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae10.103 (Index 343) (SNMP ifIndex 614) (Generation 152)
    Description: matched odd
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.103 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             3          0           194            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      ge-0/0/3.103
        Input :             0          0             0            0
        Output:             3          0           194            0
      ge-0/0/4.103
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-0/0/3.103               0           0            0            0
      ge-0/0/4.103               0           0            0            0

    List-Type       Status
    Primary         Active
                Interfaces:
                ge-0/0/4        Up
    List-Type       Status
    Backup          Waiting
                Interfaces:
                ge-0/0/3        Up
    List-Type       Status
    Standby         Down

    Protocol inet, MTU: 8978, Generation: 194, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.3.0.0/15, Local: 10.3.0.1, Broadcast: 10.3.0.3, Generation:
150
    Protocol multiservice, MTU: Unlimited, Generation: 195, Route table: 0
      Policer: Input: __default_arp_policer__

  Logical interface ae10.104 (Index 342) (SNMP ifIndex 616) (Generation 151)
    Description: matched even
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.104 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             2          0            92            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
```



```
      Adaptive Updates:          0
    Link:
      ge-0/0/3.104
        Input :             0          0             0             0
        Output:             2          0            92             0
      ge-0/0/4.104
        Input :             0          0             0             0
        Output:             0          0             0             0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-0/0/3.104                0           0            0            0
      ge-0/0/4.104                0           0            0            0

    List-Type       Status
    Primary         Active
                Interfaces:
                ge-0/0/3        Up
    List-Type       Status
    Backup          Waiting
                Interfaces:
                ge-0/0/4        Up
    List-Type       Status
    Standby         Down

    Protocol inet, MTU: 8978, Generation: 192, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.4.0.0/16, Local: 10.4.0.1, Broadcast: 10.4.0.3, Generation:
148
    Protocol multiservice, MTU: Unlimited, Generation: 193, Route table: 0
      Policer: Input: __default_arp_policer__


  Logical interface ae10.32767 (Index 341) (SNMP ifIndex 613) (Generation 150)
    Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0             0
        Output:             0          0             0             0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      ge-0/0/3.32767
        Input :             0          0             0             0
        Output:            95          0         38039             0
      ge-0/0/4.32767
        Input :             0          0             0             0
        Output:            95          0         38039             0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-0/0/3.32767              0           0            0            0
      ge-0/0/4.32767              0           0            0            0
    Protocol multiservice, MTU: Unlimited, Generation: 191, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

See Also

- distribution-list on page 1122

- targeted-distribution on page 1392
- targeted-options on page 1393

CHAPTER 8

Configuring Ethernet Automatic Protection Switching for High Availability

Ethernet Automatic Protection Switching Overview

Ethernet automatic protection switching (APS) is a linear protection scheme designed to protect VLAN based Ethernet networks.

With Ethernet APS, a protected domain is configured with two paths, a working path and a protection path. Both working and protection paths can be monitored using an Operations Administration Management (OAM) protocol like Connectivity Fault Management (CFM). Normally, traffic is carried on the working path (that is, the working path is the active path), and the protection path is disabled. If the working path fails, its protection status is marked as degraded (DG) and APS switches the traffic to the protection path, then the protection path becomes the active path.

APS uses two modes of operation, linear 1+1 protection switching architecture and linear 1:1 protection switching architecture. The linear 1+1 protection switching architecture operates with either unidirectional or bidirectional switching. The linear 1:1 protection switching architecture operates with bidirectional switching.

In the linear 1+1 protection switching architecture, the normal traffic is copied and fed to both working and protection paths with a permanent bridge at the source of the protected domain. The traffic on the working and protection transport entities is transmitted simultaneously to the sink of the protected domain, where a selection between the working and protection transport entities is made.

In the linear 1:1 protection switching architecture, the normal traffic is transported on either the working path or on the protection path using a selector bridge at the source of the protection domain. The selector at the sink of the protected domain selects the entity that carries the normal traffic.

Unidirectional and Bidirectional Switching

Unidirectional switching utilizes fully independent selectors at each end of the protected domain. Bidirectional switching attempts to configure the two end points with the same bridge and selector settings, even for a unidirectional failure. Unidirectional switching can protect two unidirectional failures in opposite directions on different entities.

Selective and Merging Selectors

In the linear 1:1 protection switching architecture, where traffic is sent only on the active path, there are two different ways in which the egress direction (the direction out of the protected segment) data forwarding can act: selective selectors and merging selectors. A selective selector forwards only traffic that is received from both the paths regardless of which one is currently active. In other words, with a merging selector the selection of the currently active path only affects the ingress direction. Merging selectors minimize the traffic loss during a protection switch, but they do not guarantee the delivery of the data packets in order.

Revertive and Nonrevertive Switching

For revertive switching, traffic is restored to the working path after the conditions causing the switch have cleared.

For nonrevertive switching, traffic is allowed to remain on the protection path even after the conditions causing the switch have cleared.

NOTE: **The configuration on both the provider edge (PE) routers have to be either in revertive mode or non-revertive mode.**

Protection Switching Between VPWS Pseudowires

Figure 11: Connections Terminating on Single PE

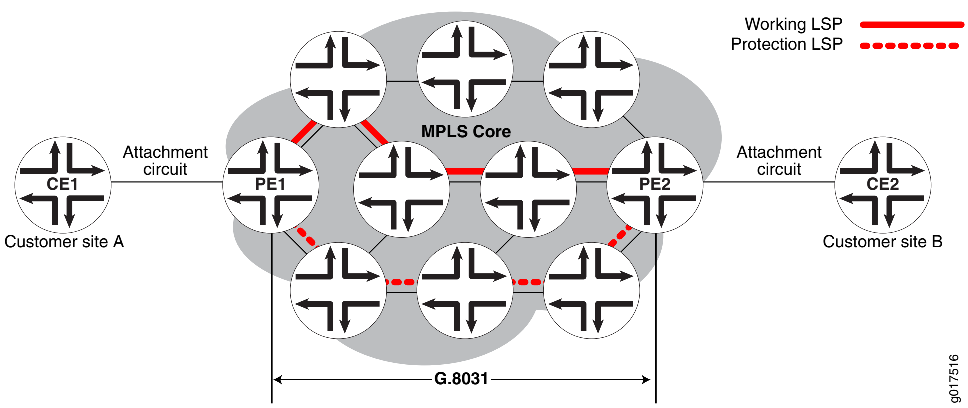

In the scenario diagramed in Figure 11 on page 214, a Virtual Private Wire Service (VPWS) is provisioned between customer sites A and B using a single pseudowire (layer 2 circuit) in the core network, and two Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) are provisioned, one for the working path and the other one for the protection path. CFM CCM will be used to monitor the status of each LSP. Provider edge routers PE1 and PE2 run G.8031 Ethernet APS to select one of the LSPs as the active path. Once the active path is elected at the source end of the protection group, PE1 forwards to traffic from site A to the elected active path. At the sink end of the protection group, PE2 implements a merging selector, meaning it forwards the traffic coming from both the LSPs to the customer site B.

Figure 12: Connections Terminating on a Different PE

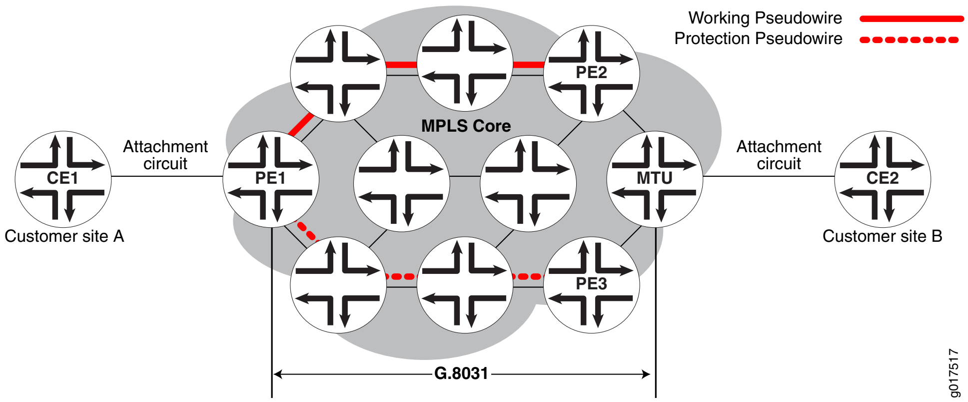

In the scenario represented in Figure 12 on page 215, a VPWS is provisioned between customer sites A and B using two pseudowires (layer 2 circuit) in the core network, one for the working path and the other for the protection path. CFM CCM will be used to monitor the status of each pseudowire.

Provider edge router PE1 and MTU run G.8031 Ethernet APS to select one of the pseudowires as the active path. Once the active path is elected at the source end of the protection group, PE1 forwards the traffic from site A to the elected active path. At the sink end of the protection group, MTU implements a merging selector, meaning it forwards the traffic coming from both the pseudowires to customer site B.

CLI Configuration Statements

```
[edit protocols protection-group]
ethernet-aps profile1{
   protocol g8031;
   revert-time seconds;
   hold-time 0-10000ms;
   local-request lockout;
}
```

revert-time- By default, protection logic restores the use of the working path once it recovers. The revert-time statement specifies how much time should elapse before the path for data should be switched from Protection to Working once recovery for Working has occurred. A revert-time of zero indicates no reversion. It will default to 300 sec (5 minutes) if not configured.

hold-time- Once a failure is detected, APS waits until this timer expires before initiating the protection switch. The range of the hold-time timer is 0 to 10,000 milliseconds. It will default to zero if not configured.

local-request- Configuring this value to lockout or force-switch will trigger lockout or force-switch operation on the protection groups using this profile.

Related Documentation

- Mapping of CCM Defects to APS Events on page 216
- Example: Configuring Protection Switching Between Psuedowires on page 217

Mapping of CCM Defects to APS Events

The continuity check message (CCM) engine marks the status of working and protected transport entities as either Down, Degraded, or Up.

Down—The monitored path is declared down if any of the following Multiple End Point (MEP) defects occur:

- Interface down
- CCM expiry
- RDI indicating signal failure

Degraded—The monitored path is declared degraded if any of the following MEP defects occur:

- FRR on
- FRR-ACK on

Up—The monitored path is declared up in the absence of any of the above events.

Figure 13: Understanding APS Events

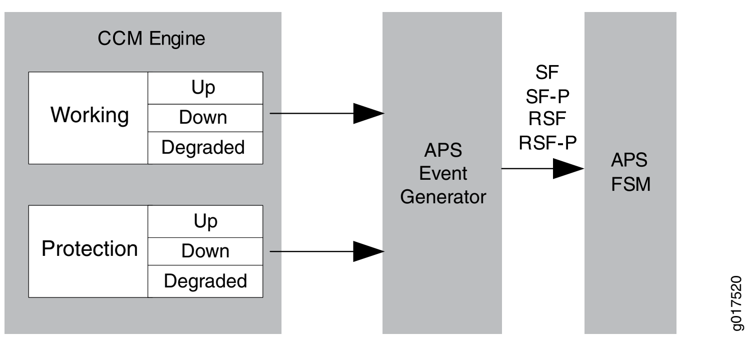

As show in Figure 13 on page 216, the APS event generator generates the following APS events based on the status of the working and protection paths:

- **SF**—Signal failure on working path
- **RSF**—Working path recovers from signal failure
- **SF-P**—Signal failure on protection path
- **RSF-P**—Protection path recovers from signal failure

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 213
- Example: Configuring Protection Switching Between Psuedowires on page 217

Example: Configuring Protection Switching Between Psuedowires

- Requirements on page 217
- Overview and Topology on page 217
- Configuration on page 217

Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.2 or later
- 2 MX Series PE routers

Overview and Topology

The physical topology of the protection switching between psuedowires example is shown in Figure 14 on page 217.

Figure 14: Topology of a Network Using VPWS Psuedowires

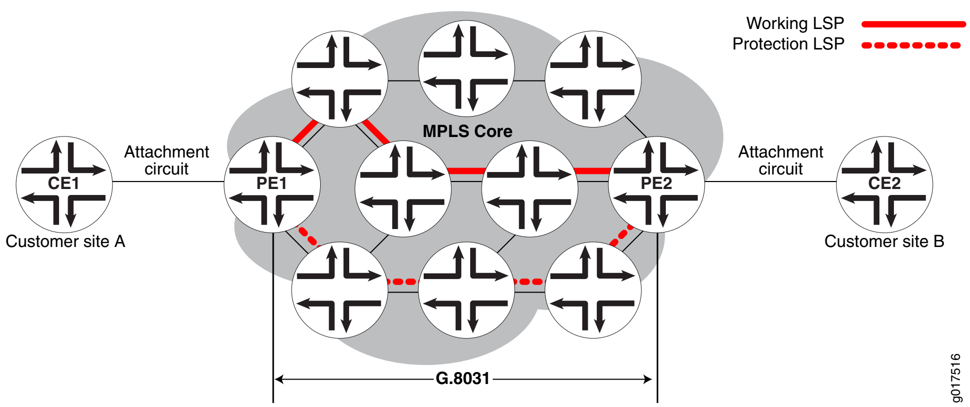

The following definitions describe the meaning of the device abbreviations used in Figure 14 on page 217.

- Customer edge (CE) device—A device at the customer site that provides access to the service provider's VPN over a data link to one or more provider edge (PE) routers.
- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.

Configuration

Step-by-Step Procedure

To configure protection switching between psuedowires, perform these tasks:

1. Configure automatic protection switching.

```
protocols {
  protection-group {
    ethernet-aps {
      profile-1 {
        protocol g8031;
        hold-time 1000s;
        revert-time 5m;
      }
    }
  }
```

2. Configure the connectivity fault management.

```
ethernet {
  oam {
    connectivity-fault-management {
      maintenance-domain md1 {
      level 5;
```

3. Configure the continuity check message for the working path.

```
maintenance-association W {
  protect maintenance-association P {
    aps-profile profile-1;
  }
  continuity-check {
    interval 1s;
  }
  mep 100 {
    interface ge-1/0/0.0 working;
    direction down;
    auto-discovery;
  }
}
```

4. Configure the continuity check message for the protection path.

```
maintenance-association P {
  continuity-check {
    interval 1s;
  }
  mep 100 {
    interface ge-1/0/0.0 protect;
    direction down;
    auto-discovery;
  }
}
```

Results Check the results of the configuration:

```
protocols {
  protection-group {
    ethernet-aps {
```



```
            profile-1 {
                protocol g8031;
                hold-time 1000s;
                revert-time 5m;
            }
        }
    }
    ethernet {
        oam {
            connectivity-fault-management {
                maintenance-domain md1 {
                    level 5;
                    maintenance-association W {
                        protect maintenance-association P {
                            aps-profile profile-1;
                        }
                        continuity-check {
                            interval 1s;
                        }
                        mep 100 {
                            interface ge-1/0/0.0 working;
                            direction down;
                            auto-discovery;
                        }
                    }
                    maintenance-association P {
                        continuity-check {
                            interval 1s;
                        }
                        mep 100 {
                            interface ge-1/0/0.0 protect;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
            }
        }
    }
```

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 213
- Mapping of CCM Defects to APS Events on page 216

CHAPTER 9

Configuring Ethernet Ring Protection Switching for High Availability

Ethernet Ring Protection Switching Overview

Ethernet ring protection switching (ERPS) helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. The RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An automatic protection switching (APS) protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.

NOTE: ERPS on AE interfaces is not supported on ACX Series routers.

The following standards provide detailed information on Ethernet ring protection switching:

- ITU-T Recommendation G.8032/Y.1344 version 1 and 2, *Ethernet Ring protection switching*. G.8032v1 supports a single ring topology and G.8032v2 supports multiple rings and ladder topology.

  NOTE: EX2300 and EX3400 switches support G.8032v1 only.

- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet-based networks*

For additional information on configuring Ethernet ring protection switching on EX Series switches, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*.

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *Layer 2 Configuration Guide* for a complete example of Ethernet rings and information about STP loop avoidance and prevention.

Related Documentation

- Understanding Ethernet Ring Protection Switching Functionality on page 222
- Configuring Ethernet Ring Protection Switching on page 229
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 230
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding Ethernet Ring Protection Switching Functionality

Acronyms

The following acronyms are used in the discussion about Ethernet ring protection switching (ERPS):

- MA—Maintenance association
- MEP—Maintenance association end point
- OAM—Operations, administration, and management (Ethernet ring protection switching uses connectivity fault management daemon)
- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—Ring automatic protection switching
- WTB—Wait to block. Note that WTB is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting on EX2300 and EX3400 switches has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect on EX2300 and EX4300 switches.
- WTR—Wait to restore. Note that on EX2300 and EX3400 switches only, the WTR configuration must be 5-12 minutes.
- RPL—Ring protection link

Ring Nodes

Multiple nodes are used to form a ring. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL.

Ring Node States

The following are the different states for each node of a specific ring:

- init—Not a participant of a specific ring.
- idle—No failure on the ring; the node is performing normally. For a normal node, traffic is unblocked on both ring ports. For the RPL owner or RPL neighbor, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- protection—A failure occurred on the ring. For a normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.
- pending—The node is recovering from failure or its state after a **clear** command is used to remove the previous manual command. When a protection group is configured, the node enters the pending state. When a node is in pending state, the WTR or WTB timer will be running. All nodes are in pending state till WTR or WTB timer expiry.

- force switch—A force switch is issued. When a force switch is issued on a node in the ring all nodes in the ring will move into the force switch state.

NOTE: **EX2300 and EX3400 switches do not support force switch.**

- manual switch—A manual switch is issued. When a manual switch is issued on a node in the ring all nodes in the ring will move into the manual switch state.

NOTE: **EX2300 and EX3400 switches do not support manual switch.**

There can be only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

Default Logging of Basic State Transitions on EX Series Switches

Starting with Junos OS Release 14.1X53-D15, EX Series switches automatically log basic state transitions for the ERPS protocol. Starting with Junos OS Release 18.2R1, EX2300 and EX3400 switches automatically log basic state transitions for the ERPS protocol. No configuration is required to initiate this logging. Basic state transitions include ERPS interface transitions from up to down, and down to up; and ERPS state transitions from idle to protection, and protection to idle.

The basic state transitions are logged in a single file named **erp-default**, which resides in the **/var/log** directory of the switch. The maximum size of this file is 15 MB.

Default logging for ERPS can capture initial ERPS interface and state transitions, which can help you troubleshoot issues that occur early in the ERPS protocol startup process. However, if more robust logging is needed, you can enable traceoptions for ERPS by entering the **traceoptions** statement in the **[edit protocols protection-group]** hierarchy.

Be aware that for ERPS, only default logging or traceoptions can be active at a time on the switch. That is, default logging for ERPS is automatically enabled and if you enable traceoptions for ERPS, the switch automatically disables default logging. Conversely, if you disable traceoptions for ERPS, the switch automatically enables default logging.

Logical Ring

You can define multiple logical-ring instances on the same physical ring. The logical ring feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring must operate on the physical interface, not the VLAN. Multiple ring instances are usually defined with trunk mode ring interfaces.

FDB Flush

When ring protection switching occurs, normally an *FDB flush* is executed. The Ethernet ring control module uses the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.

NOTE: Optimized flushing is not supported on EX2300 and EX3400 switches.

Starting with Junos OS Release 14.2, the FDB flush depends on the RAPS messages received on the both the ports of the ring node.

Traffic Blocking and Forwarding

Ethernet ring control uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical interface default STP index state to forwarding or discarding in order to control user traffic.

RPL Neighbor Node

Starting with Junos OS Release 14.2, ring protection link neighbor nodes are supported. An RPL neighbor node is adjacent to the RPL and is not the RPL owner. If a node is configured with one interface as the protection-link-end and no protection-link-owner is present in its configuration, the node is an RPL neighbor node.

NOTE: RPL neighbor node is not supported on EX2300 and EX3400 switches.

RAPS Message Blocking and Forwarding

The router or switch treats the ring automatic protection switching (RAPS) message the same as it treats user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS messages between the two ring ports, as shown in Figure 15 on page 225, the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces' STP index state. The RAPS message is always sent by the router or switch through the ring ports, as shown in Figure 16 on page 225. A RAPS message received from a discarding ring port is sent to the Ethernet ring control module, but is not sent to the other ring port.

Figure 15: Protocol Packets from the Network to the Router

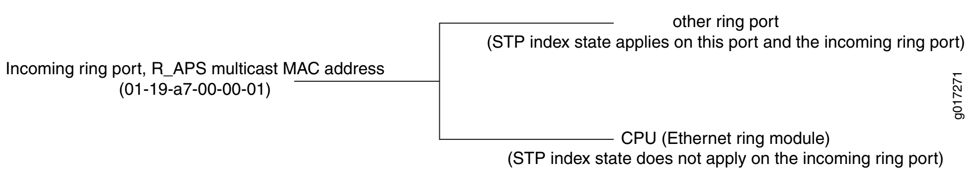

Figure 16: Protocol Packets from the Router or Switch to the Network

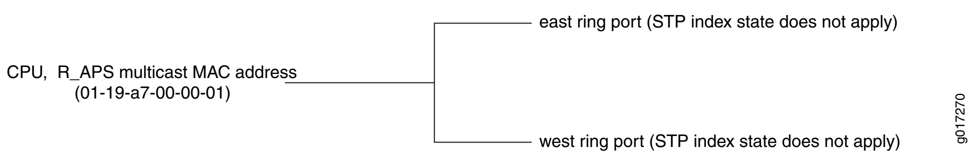

Juniper Networks switches and Juniper Networks routers use different methods to achieve these routes.

The switches use forwarding database entries to direct the RAPS messages. The forwarding database entry (keyed by the RAPS multicast address and VLAN) has a composite next hop associated with it—the composite next hop associates the two ring interfaces with the forwarding database entry and uses the split horizon feature to prevent sending the packet out on the interface that it is received on. This is an example of the forwarding database entry relating to the RAPS multicast MAC (a result of the **show ethernet-switching table detail** command):

```
VLAN: v1, Tag: 101, MAC: 01:19:a7:00:00:01, Interface: ERP
Interfaces:                  ge-0/0/9.0, ge-0/0/3.0
Type: Static
Action: Mirror
Nexthop index: 1333
```

The routers use an implicit filter to achieve ERP routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:
  - ```
    term 1: if [Ethernet type is not OAM Ethernet type (0x8902)
    ]               { accept packet }
    ```
  - ```
    term 2: if [source MAC address belongs to this bridge]
                 { drop packet, our packet loop through the ring and come back
                     to home}
    ```
  - ```
    term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,
                      0x00,0x00,0x01] AND[ring port STP status is DISCARDING]
                  { send to CPU }
    ```
- Control channel related terms:
  - ```
          if [destination is the RAPS PDU multicast
    address(0x01,0x19,0xa7,0x00,0x00,
             0x01] AND[ring port STP status is FORWARDING] AND [Incoming
    interface
              IFL equal to control channel IFL]
             { send packet to CPU and send to the other ring port }
          default term: accept packet.
    ```

Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain for routers (or the same VLAN for switches) in order to forward RAPS protocol data units (PDUs) between the two ring control physical interfaces. If the router control channel logical interface is not a trunk port, only control logical interfaces will be configured in ring port configuration. If this router control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured for routers. For switches, always specify either a VLAN name or VLAN ID for all links.

RAPS Message Termination

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

Revertive and Non-revertive Modes

In revertive operation, once the condition causing a switch has cleared, traffic is blocked on the RPL and restored to the working transport entity. In nonrevertive operation, traffic is allowed to use the RPL if it has not failed, even after a switch condition has cleared.

NOTE: Non-revertive mode is not supported on EX2300 and EX3400 switches.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). The ring control module also supports the interconnection of multiple rings. Interconnection of two rings means that two rings might share the same link or share the same node. Ring interconnection is supported only using non-virtual-channel mode. Ring interconnection using virtual channel mode is not supported.

NOTE: Interconnection of multiple rings is not supported on EX2300 and EX3400 switches.

Node ID

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address.

For routers only, you can configure this node ID when configuring the ring on the node or automatically select an ID like STP does. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID. The node ID on switches is selected automatically and is not configurable.

Ring ID

The ring ID is used to determine the value of the last octet of the MAC destination address field of the RAPS protocol data units (PDUs) generated by the ERP control process. The ring ID is also used to discard any RAPS PDU, received by this ERP control process with a non-matching ring ID. Ring ID values 1 through 239 are supported.

Bridge Domains with the Ring Port (MX Series Routers Only)

On the routers, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain, you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB Layer 3 interface.

Wait-to-Block Timer

The RPL owner node uses a delay timer before initiating an RPL block in revertive mode of operation or before reverting to IDLE state after clearing manual commands. The Wait-to-Block (WTB) timer is used when clearing **force switch** and **manual switch** commands. As multiple **force switch** commands are allowed to coexist in an Ethernet ring, the WTB timer ensures that clearing of a single **force switch** command does not trigger the re-blocking of the RPL. When clearing a **manual switch** command, the WTB timer prevents the formation of a closed loop due to a possible timing anomaly where the RPL Owner Node receives an outdated remote **manual switch** request during the recovery process.

When recovering from a **manual switch** command, the delay timer must be long enough to receive any latent remote **force switch**, signal failure, or **manual switch** commands. This delay timer is called the WTB timer and is defined to be 5 seconds longer than the guard timer. This delay timer is activated on the RPL Owner Node. When the WTB timer expires, the RPL Owner Node initiates the reversion process by transmitting an RAPS (NR, RB) message. The WTB timer is deactivated when any higher-priority request preempts it.

NOTE: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.

Adding and Removing a Node

Starting with Junos OS Release 14.2, you can add or remove a node between two nodes in an Ethernet ring. Nodes are added or removed using the **force switch** command.

NOTE: EX2300 and EX3400 switches do not support force switch.

Release History Table

Release	Description
18.2R1	Starting with Junos OS Release 18.2R1, EX2300 and EX3400 switches automatically log basic state transitions for the ERPS protocol.
14.2	Starting with Junos OS Release 14.2, the FDB flush depends on the RAPS messages received on the both the ports of the ring node.
14.2	Starting with Junos OS Release 14.2, ring protection link neighbor nodes are supported.
14.2	Starting with Junos OS Release 14.2, you can add or remove a node between two nodes in an Ethernet ring.
14.1X53-D15	Starting with Junos OS Release 14.1X53-D15, EX Series switches automatically log basic state transitions for the ERPS protocol.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Configuring Ethernet Ring Protection Switching on page 229
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 230

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

Configuring Ethernet Ring Protection Switching

The inheritance model follows:

```
protection-group {
  ethernet-ring ring-name (
    node-id mac-address;
    ring-protection-link-owner;
    east-interface {
      control-channel channel-name {
      ring-protection-link-end;
    }
    west-interface {
      node-id mac-address;
      control-channel channel-name {
      ring-protection-link-end;
    }
    data-channel {
      vlan number;
    }
    guard-interval number;
    restore-interval number;
  }
}
```

For each ring, a protection group must be configured. There may be several rings in each node, so there should be multiple protection groups corresponding to the related Ethernet rings.

Three interval parameters (**restore-interval**, **guard-interval**, and **hold-interval**) can be configured at the protection group level. These configurations are global configurations and apply to all Ethernet rings if the Ethernet ring doesn't have a more specific configuration for these values. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Understanding Ethernet Ring Protection Switching Functionality on page 222
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 230
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Ethernet Ring Protection Switching Configuration on MX Routers

This example describes how to configure Ethernet ring protection switching on an MX Series router:

Requirements

This example uses the following hardware and software components:

- Router node 1 running Junos OS with two Gigabit Ethernet interfaces.
- Router node 2 running Junos OS with two Gigabit Ethernet interfaces.
- Router node 3 running Junos OS with two Gigabit Ethernet interfaces.

Ethernet Ring Overview and Topology

This section describes a configuration example for a three-node ring. The ring topology is shown in Figure 17 on page 231.

Figure 17: Example of a Three-Node Ring Topology

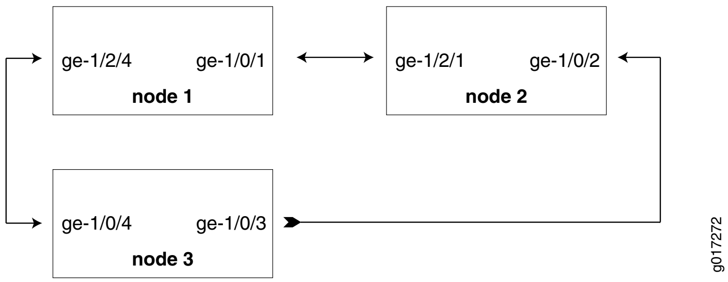

The configuration in this section is only for the RAPS channel. The bridge domain for user traffic is the same as the normal bridge domain. The only exception is if a bridge domain includes a ring port, then it must also include the other ring port of the same ring.

Configuring a Three-Node Ring

To configure Ethernet Ring Protection Switching on a three-node ring, perform these tasks:

- Configuring Ethernet Ring Protection Switching on a Three-Node Ring on page 231

Configuring Ethernet Ring Protection Switching on a Three-Node Ring

Step-by-Step Procedure

1. **Configuring Node 1**

```
interfaces {
     ge-1/0/1 {
          vlan-tagging;
          encapsulation flexible-ethernet-services;
          unit 1 {
               encapsulation vlan-bridge;
               vlan-id 100;
          }
     }
     ge-1/2/4 {
          vlan-tagging;
          encapsulation flexible-ethernet-services;
          unit 1 {
               encapsulation vlan-bridge;
               vlan-id 100;
               }
          }
     }
     bridge-domains {
          bd1 {
               domain-type bridge;
               interface ge-1/2/4.1;
               interface ge-1/0/1.1;
               }
          }
     protocols {
          protection-group {
               ethernet-ring pg101 {
                    node-id 00:01:01:00:00:01;
                    ring-protection-link-owner;
                    east-interface {
                         control-channel ge-1/0/1.1;
```

```
                    ring-protection-link-end;
                    }
                west-interface {
                    control-channel ge-1/2/4.1;
                    }
                }
            }
        }
    protocols {
        oam {
            ethernet {
                connectivity-fault-management {
                    action-profile rmep-defaults {
                        default-action {
                            interface-down;
                            }
                        }
                    maintenance-domain d1 {
                        level 0;
                        maintenance-association 100 {
                            mep 1 {
                                interface ge-1/0/1;
                                remote-mep 2 {
                                   action-profile rmep-defaults;

                                    }
                                }
                            }
                        }
                    maintenance-domain d2 {
                        level 0;
                        maintenance-association 100 {
                            mep 1 {
                                interface ge-1/2/4;
                                remote-mep 2 {
                                   action-profile rmep-defaults;

                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
    }
```

2. **Configuring Node 2**

```
interfaces {
    ge-1/0/2 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
             encapsulation vlan-bridge;
             vlan-id 100;
        }
    }

    ge-1/2/1 {
        vlan-tagging;
```



```
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
}

bridge-domains {
    bd1 {
        domain-type bridge;
        interface ge-1/2/1.1;
        interface ge-1/0/2.1;
    }
}


protocols {
    protection-group {
        ethernet-ring pg102 {
            east-interface {
                control-channel ge-1/0/2.1;
            }
            west-interface {
                control-channel ge-1/2/1.1;
            }
        }
    }
}


protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
                maintenance-domain d1 {
                    level 0;
                    maintenance-association 100 {
                        mep 2 {
                            interface ge-1/2/1;
                            remote-mep 1 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
                maintenance-domain d3 {
                    level 0;
                    maintenance-association 100 {
                        mep 1 {
                            interface ge-1/0/2;
                            remote-mep 2 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
```

```
                }
            }
        }
    }
}
```

3. **Configuring Node 3**

```
interfaces {
    ge-1/0/4 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
             encapsulation vlan-bridge;
             vlan-id 100;
        }
    }

    ge-1/0/3 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
}

bridge-domains {
    bd1 {
        domain-type bridge;
        interface ge-1/0/4.1;
        interface ge-1/0/3.1;
    }
}


protocols {
    protection-group {
        ethernet-ring pg103 {
            east-interface {
                control-channel ge-1/0/3.1;
            }
            west-interface {
                control-channel ge-1/0/4.1;
            }
        }
    }
}


protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
                maintenance-domain d2 {
```



```
                        level 0;
                        maintenance-association 100 {
                            mep 2 {
                                interface ge-1/0/4;
                                remote-mep 1 {
                                    action-profile rmep-defaults;
                                }
                            }
                        }
                    }
                    maintenance-domain d3 {
                        level 0;
                        maintenance-association 100 {
                            mep 2 {
                                interface ge-1/0/3;
                                remote-mep 1 {
                                    action-profile rmep-defaults;
                                }
                            }
                        }
                    }
                }
            }
        }
    }
```

Examples: Ethernet RPS Output

This section provides output examples based on the configuration shown in "Example: Ethernet Ring Protection Switching Configuration on MX Routers" on page 230. The show commands used in these examples can help verify configuration and correct operation.

Normal Situation—RPL Owner Node

If the ring has no failure, the **show** command will have the following output for Node 1:

```
user@node1> show protection-group ethernet-ring aps

Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               NR             No        Yes

Originator  Remote Node ID
Yes

user@node1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/0/1     ge-1/0/1.1       discarding     Yes
ge-1/2/4     ge-1/2/4.1       forwarding     No

Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready

user@node1> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         Ring Protection Link Owner
pg101            idle         NR-RB         Yes

Restore Timer  Quard Timer  Operation state
disabled       disabled    operational
```

```
user@node1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                          : 1
RAPS received                      : 0
Local SF happened:                 : 0
Remote SF happened:                : 0
NR event happened:                 : 0
NR-RB event happened:              : 1
```

Normal Situation—Other Nodes

For Node 2 and Node 3, the outputs should be the same:

```
user@node2> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               NR             No        Yes

Originator  Remote Node ID
No          00:01:01:00:00:01
```

```
user@node2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/2/1     ge-1/2/1.1       forwarding     No
ge-1/0/2     ge-1/0/2.1       forwarding     No

Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```

```
user@node2> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         Ring Protection Link Owner
pg102            idle         NR-RB         No

Restore Timer  Quard Timer  Operation state
disabled       disabled    operational
```

```
user@node2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg101
RAPS sent                          : 0
RAPS received                      : 1
Local SF happened:                 : 0
Remote SF happened:                : 0
NR event happened:                 : 0
NR-RB event happened:              : 1
```

Failure Situation—RPL Owner Node

If the ring has a link failure between Node 2 and Node 3, the **show** command will have the following outputs for Node 1:

```
user@node1> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               SF             NO        No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

```
user@node1> show protection-group ethernet-ring interface
```



```
Ethernet ring port parameters for protection group pg101

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/0/1     ge-1/0/1.1       forwarding     Yes
ge-1/2/4     ge-1/2/4.1       forwarding     No

Signal Failure  Admin State
Clear           IFF ready
Clear           IFF ready
```

```
user@node1> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         Ring Protection Link Owner
pg101            protected    SF            Yes

Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

```
user@node1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                        : 1
RAPS received                    : 1
Local SF happened:               : 0
Remote SF happened:              : 1
NR event happened:               : 0
NR-RB event happened:            : 1
```

Failure Situation—Other Nodes

For Node 2 and Node 3, the outputs should be the same:

```
user@node2> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               SF             No        No

Originator  Remote Node ID
Yes         00:00:00:00:00:00
```

```
user@node2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/2/1     ge-1/2/1.1       forwarding     No
ge-1/0/2     ge-1/0/2.1       discarding     No

Signal Failure  Admin State
Clear           IFF ready
set             IFF ready
```

```
user@node2> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         Ring Protection Link Owner
pg102            idle         NR-RB         No

Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

```
user@node2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg101
RAPS sent                        : 1
RAPS received                    : 1
Local SF happened:               : 1
Remote SF happened:              : 0
NR event happened:               : 0
NR-RB event happened:            : 1
```

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Understanding Ethernet Ring Protection Switching Functionality on page 222
- Configuring Ethernet Ring Protection Switching on page 229
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 10

Configuring MAC Address Validation on Static Ethernet Interfaces

MAC Address Validation on Static Ethernet Interfaces Overview

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

MAC address validation is supported on AE, Fast Ethernet, Gigabit Ethernet, and 10–Gigabit Ethernet interfaces (with or without VLAN tagging) on MX Series routers only.

There are two types of MAC address validation that you can configure:

- Loose—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

  Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not support the MAC address of the tuple

  Continues to forward packets when the source address of the incoming packet does not match any of the trusted IP addresses.

- Strict—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

  Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Related Documentation

- Configuring MAC Address Validation on Static Ethernet Interfaces on page 240
- Disabling MAC Address Learning of Neighbors Through ARP or Neighbor Discovery for IPv4 and IPv6 Neighbors on page 241

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring MAC Address Validation on Static Ethernet Interfaces

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address. MAC address validation is supported on AE, Fast Ethernet, Gigabit Ethernet, and 10–Gigabit Ethernet interfaces (with or without VLAN tagging) on MX Series routers only.

To configure MAC address validation on static Ethernet Interfaces:

1. In the configuration mode, at the [**edit**] hierarchy level, configure the static Ethernet interface.

   [edit]
   user@host# edit interfaces *interface-name*

2. Configure the protocol family and the logical unit of the interface at the [**edit interfaces *interface-name***] hierarchy level. While configuring the protocol family, specify **inet** as the protocol family.

   [edit interfaces *interface-name*]
   user@host# edit unit *logical-unit-number* family inet

3. Configure MAC address validation on the static Ethernet Interface. You can specify the type of MAC address validation you require. Possible values are: Strict and Loose. You can also specify the interface address.

   [edit interfaces *interface-name*] unit *logical-unit-number* family inet
   user@host# set mac-validate *option* address *address*

4. Configure the static ARP entry by specifying the IP address and the MAC address that are to be mapped. The IP address specified must be part of the subnet defined in the enclosing **address** statement. The MAC address must be specified as hexadecimal bytes in the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn*** format. For instance, you can use either **0011.2233.4455** or **00:11:22:33:44:55**.

   [edit interfaces *interface-name* unit *logical-unit-number* family inet address
     interface-address
   user@host# set arp *ip-address* mac *mac-address*

Related Documentation

- family on page 1162
- mac-validate on page 1247
- MAC Address Validation on Static Ethernet Interfaces Overview on page 239

- *Ethernet Interfaces Feature Guide for Routing Devices*

Disabling MAC Address Learning of Neighbors Through ARP or Neighbor Discovery for IPv4 and IPv6 Neighbors

The Junos OS provides the **no-neighbor-learn** configuration statement at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** and **[edit interfaces *interface-name* unit *interface-unit-number* family inet6]** hierarchy levels.

To disable ARP address learning by not sending arp-requests and not learning from ARP replies for IPv4 neigbors, include the **no-neighbor-learn** statement at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet**] hierarchy level:

```
[edit interfaces interface-name unit interface-unit-number family inet]
no-neighbor-learn;
```

To disable neighbor discovery for IPv6 neighbors, include the **no-neighbor-learn** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level:

```
[edit interfaces interface-name unit interface-unit-number family inet6]
no-neighbor-learn;
```

Related Documentation

- *Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 11

Configuring 802.1Q VLANs

802.1Q VLANs Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Related Documentation

- Configuring Dynamic 802.1Q VLANs on page 246
- 802.1Q VLAN IDs and Ethernet Interface Types on page 245
- Enabling VLAN Tagging on page 247
- Binding VLAN IDs to Logical Interfaces on page 251
- Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs on page 270
- Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 257
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 257
- Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 258
- Specifying the Interface to Handle Traffic for a CCC on page 258
- Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 267
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 266
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 268
- Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 259
- Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 268
- Configuring Access Mode on a Logical Interface on page 261
- Configuring a Logical Interface for Trunk Mode on page 262
- Configuring the VLAN ID List for a Trunk Interface on page 262
- Configuring a Trunk Interface on a Bridge Network on page 263
- *Ethernet Interfaces Feature Guide for Routing Devices*

802.1Q VLAN IDs and Ethernet Interface Types

You can partition the router into up to 4095 different VLANs—depending on the router model and the physical interface types—by associating logical interfaces with specific VLAN IDs.

VLAN ID 0 is reserved for tagging the priority of frames. VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN circuit cross-connect (CCCs).

For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. With flexible Ethernet services encapsulation, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

The maximum number of user-configurable VLANs is 15 on each port of the Dense-FE PIC (8-port/12-port/48-port).

Table 14 on page 245 lists VLAN ID range by interface type.

Table 14: VLAN ID Range by Interface Type

Interface Type	VLAN ID Range
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
4-port, 8-port, and 12-port Fast Ethernet	1 through 1023
48-port Fast Ethernet	1 through 4094
Tri-Rate Ethernet copper	1 through 4094
Gigabit Ethernet	1 through 4094
Gigabit Ethernet IQ	1 through 4094
10-Gigabit Ethernet	1 through 4094
100-Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023

NOTE: For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the built-in Gigabit Ethernet port on the M7i router), VLAN IDs on a single interface can differ from each other.

Because IS-IS has an 8-bit limit for broadcast multiaccess media, you cannot set up more than 255 adjacencies over Gigabit Ethernet using VLAN tagging. For more information, see the *Junos OS Routing Protocols Library*.

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Dynamic 802.1Q VLANs

You can configure the router to dynamically create VLANs when a client accesses an interface and requests a VLAN ID that does not yet exist. When a client accesses a VLAN interface, the router instantiates a VLAN dynamic profile that you have associated with the interface. Using the settings in the dynamic profile, the router extracts information about the client from the incoming packet (for example, the interface and unit values), saves this information in the routing table, and creates a VLAN or stacked VLAN ID for the client from a range of VLAN IDs that you configure for the interface.

Dynamically configuring VLANs or stacked VLANs requires the following general steps:

1. Configure a dynamic profile for dynamic VLAN or dynamic stacked VLAN creation.
2. Associate the VLAN or stacked VLAN dynamic profile with the interface.
3. Specify the Ethernet packet type that the VLAN dynamic profile accepts.
4. Define VLAN ranges for use by the dynamic profile when creating VLAN IDs.

For procedures on how to configure dynamic VLANs and dynamic stacked VLANs for client access, see the *Junos OS Broadband Subscriber Management and Services Library*.

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling VLAN Tagging

You can configure the router to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

NOTE: **If you configure VLAN tagging on Gigabit Ethernet IQ, IQ2, and IQ2-E interfaces on M320, M120, and T Series routers, Junos OS creates an internal logical interface that reserves 50 Kbps of bandwidth from Gigabit Ethernet IQ interfaces and 2 Mbps of bandwidth from Gigabit Ethernet IQ2 and IQ2-E interfaces. As a result, the effective available bandwidth for these interface types is now 999.5 Mbps and 998 Mbps, respectively.**

1. To configure the router to receive and forward single-tag frames with 802.1Q VLAN tags, include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   [edit interfaces *interface-name*]
   user@host# vlan-tagging;

2. To configure the router to receive and forward dual-tag frames with 802.1Q VLAN tags, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   [edit interfaces *interface-name*]
   user@host# stacked-vlan-tagging;

3. Mixed tagging is supported for Gigabit Ethernet interfaces on Gigabit Ethernet IQ2 and IQ2-E, and IQ or IQE PICs on M Series and T Series routers, for all router Gigabit and 10-Gigabit Ethernet interfaces on MX Series routers, and for aggregated Ethernet interfaces with member links in IQ2 and IQ2-E PICs or in MX Series DPCs. Mixed tagging enables to configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

   NOTE: **Mixed tagging is not supported on Fast Ethernet interfaces.**

   To configure mixed tagging:

   a. Configure the **flexible-vlan-tagging** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level.

      [edit interfaces ge-*fpc/pic/port*]
      user@host#flexible-vlan-tagging;

   b. Configure the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

      [edit interfaces ge-*fpc/pic/port* unit *logival-unit-number*]

```
user@host# vlan-id number;
family family {
  address address;
}
user@host# vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
  address address;
}
```

NOTE: If you configure the physical interface MTU for mixed tagging, then you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

If the same physical interface MTU value is configured on both the VLAN and flexible VLAN-tag routers, the L2 circuit configuration does not come up and a MTU mismatch is logged. However, normal traffic flow is unaffected.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.

4. For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs, for all MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces configured for 802.1Q flexible VLAN tagging, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, you can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the **native-vlan-id** statement and the **flexible-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit interfaces ge-fpc/pic/port]
   flexible-vlan-tagging;
   native-vlan-id number;
   ```

   The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Related Documentation

- 802.1Q VLANs Overview on page 244
- Configuring VLAN and Extended VLAN Encapsulation on page 256
- Stacking a VLAN Tag on page 569

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Sending Untagged Traffic Without VLAN ID to Remote End on page 249

Sending Untagged Traffic Without VLAN ID to Remote End

Send traffic without the native VLAN ID (*native-vlan-id*) to the remote end of the network if untagged traffic is received.

If this option is not configured, then *native-vlan-id* is added to untagged traffic. But if this option is configured, then *native-vlan-id* is not added to untagged traffic.

NOTE:

- This feature works only on MX series routers with MPCs/MICs. Configuring this option with DPC results in no behavior change. But, if this option is configured with Aggregated Ethernet (AE) in which the sub interfaces reside across MPCs/MICs and DPC, MPCs/MICs and DPC will show a different behavior.
- In the egress direction, this feature is disrupted by VLAN normalization. Because of normalization, the egress interface cannot distinguish between untagged traffic and tagged traffic. And untagged traffic is sent out with *native-vlan-id*. Consider this while configuring both VLAN normalization and new *native-vlan-id* option.

  There will be a problem with ingress firewall filter if filter term includes *native-vlan-id*. With **no-native-vlan-insert** option configured, *native-vlan-id* will not be inserted to untagged traffic. So, firewall filter term will not match with untagged traffic. But if incoming traffic have VLAN ID which is equal to *native-vlan-id*, then firewall filter term will match and firewall will work.
- When this feature is used with AE, all sub-interfaces of AE should be in same type of FPC.

Related Documentation

- *Configuring Interface Encapsulation on Physical Interfaces*
- 802.1Q VLANs Overview on page 244
- *Configuring VPLS Interface Encapsulation*
- *native-vlan-id*
- **no-native-vlan-insert on page 1277**
- Enabling VLAN Tagging on page 247

Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers

This topic describes how to configure flexible VLAN tagging on PTX Series Packet Transport Routers. In addition to VLAN tagging and stacked VLAN tagging, you can configure a port for flexible tagging. With flexible VLAN tagging, you can configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces et-*fpc/pic/port*]** hierarchy level. You must also include the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces et-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces et-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
   vlan-id number;
}
unit logical-unit-number {
   vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
}
```

Related Documentation

- Enabling VLAN Tagging on page 247

Binding VLAN IDs to Logical Interfaces

This topic describes how to configure logical interfaces to receive and forward VLAN-tagged frames:

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, a range of VLAN IDs, or a list of VLAN IDs to the logical interface. Table 15 on page 251 lists the configuration statements you use to bind VLAN IDs to logical interfaces, organized by scope of the VLAN IDs used to match incoming packets. You can configure these statements at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Table 15: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces

Scope of VLAN ID Matching	Type of VLAN Framing Supported on the Logical Interface	
	Single-Tag Framing	Dual-Tag Framing
VLAN ID	**vlan-id** *vlan-id*;	**vlan-tags outer** *tpid*.<*vlan-id*> **inner** *tpidvlan-id*;
VLAN ID Range	**vlan-id-range** *vlan-id–vlan-id*;	**vlan-tags outer** *tpid.vlan-id* **inner-range** *tpid.vlan-id–vlan-id*;
VLAN ID List	**vlan-id-list** [*vlan-id vlan-id–vlan-id*];	**vlan-tags outer** <*tpid.*>*vlan-id* **inner-list** [*vlan-id vlan-id–vlan-id*];

NOTE: The **inner-list** option of the **vlan-tags** statement does not support Tag Protocol ID (TPID) values.

1. A logical interface that you have associated (bound) to a particular VLAN ID will receive and forward incoming frames that contain a matching VLAN ID. To bind a VLAN ID to a single-tag logical interface, include the **vlan-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

   [edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# vlan-id *vlan-id*;

   To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

2. To bind a VLAN ID to a dual-tag logical interface, include the **vlan-tags** statement at the **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level:

   [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
   vlan-tags inner <*tpid.*>*vlan-id* outer <*tpid.*>*vlan-id*;

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

3. A VLAN range can be used by service providers to interconnect multiple VLANs belonging to a particular customer over multiple sites. Using a VLAN ID range conserves switch resources and simplifies configuration. To bind a range of VLAN IDs to a single-tag logical interface, include the **vlan-id-range** statement at the **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level.

   [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
   vlan-id-range *vlan-id-vlan-id*;

4. To bind a range of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner–list** option to specify the VLAN IDs as an inclusive range by separating the starting VLAN ID and ending VLAN ID with a hyphen. You can include the statement at the **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level.

   [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
   vlan-tags inner–list [*vlan-id vlan-id–vlan-id*] outer <*tpid*.>*vlan-id*;

   To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

In Junos OS Release 9.5 and later, on MX Series routers and in Junos OS Release 12.2R2 and later on EX Series switches, you can bind a list of VLAN IDs to a single logical interface, eliminating the need to configure a separate logical interface for every VLAN or VLAN range. A logical interface that accepts packets tagged with any VLAN ID specified in a VLAN ID list is called a *VLAN-bundled* logical interface.

You can use VLAN-bundled logical interfaces to configure circuit cross-connects between Layer 2 VPN routing instances or Layer 2 circuits. Using VLAN-bundled logical interfaces simplifies configuration and reduces use of system resources such as logical interfaces, next hops, and circuits.

As an alternative to configuring multiple logical interfaces (one for each VLAN ID and one for each range of VLAN IDs), you can configure a single VLAN-bundled logical interface based on a list of VLAN IDs.

NOTE: The **vlan-id** option is not supported to achieve VLAN normalization on VPLS instances that are configured with **vlan-id-list**. However, you can use the **vlan-maps** option to achieve VLAN normalization.

1. To bind a list of VLAN IDs to a single-tag logical interface, include the **vlan-id-list** statement at the **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level. Specify the VLAN IDs in the list individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

   [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
   user@host# vlan-id-list [*vlan-id vlan-id–vlan-id*];

   To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

2. To bind a list of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement at the **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]** hierarchy level. Use the **inner-list** option to specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

   [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
   user@host# vlan-tags inner-list [*vlan-id vlan-id–vlan-id*] outer <*tpid*>*vlan-id*;

   NOTE: The **inner-list** option of the **vlan-tags** statement does not support Tag Protocol ID (TPID) values.

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

The following sample configuration configures two different lists of VLAN IDs on two different logical ports.

```
[edit interfaces]
ge-1/1/0 {
    vlan-tagging; # Only for single-tagging
    encapsulation flexible-ethernet-services;
    unit 10 {
        encapsulation vlan-ccc;
        vlan-id-list [20 30–40 45];
    }
}
ge-1/1/1 {
    flexible-vlan-tagging; # Only for mixed tagging
    encapsulation flexible-ethernet-services;
    unit 10 {
        encapsulation vlan-ccc;
        vlan-id-list [1 10 20 30–40];
    }
    unit 20 {
        encapsulation vlan-ccc;
        vlan-tags outer 200 inner-list [50–60 80 90–100];
    }
}
```

In the example configuration above, **ge-1/1/0** supports single-tag logical interfaces, and **ge-1/1/1** supports mixed tagging. The single-tag logical interfaces **ge-1/1/0.10** and **ge-1/1/1.20** each bundle lists of VLAN IDs. The dual-tag logical interface **ge-1/1/1.20** bundles lists of inner VLAN IDs.

TIP: You can group a range of identical interfaces into an interface range and then apply a common configuration to that interface range. For example, in the above example configuration, both interfaces ge-1/1/0 and ge-1/1/1 have the same physical encapsulation type of **flexible-ethernet-services**. Thus you can define an interface range with the interfaces ge-1/1/0 and ge-1/1/1 as its members and apply the encapsulation type flexible-ethernet-services to that defined interface range.

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Configuring Interface Ranges*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Associating VLAN IDs to VLAN Demux Interfaces

The following sections describe how to configure VLAN demux interfaces to receive and forward VLAN-tagged frames:

Associating VLAN IDs to VLAN Demux Interfaces Overview

To configure a VLAN demux interface to receive and forward VLAN-tagged frames, you must associate a VLAN ID or dual tagged (stacked) VLAN ID to the interface. Table 16 on page 255 shows the configuration statements you use to associate VLAN IDs to VLAN demux interfaces, depending on the VLAN tag framing you use:

Table 16: Configuration Statements Used to Associate VLAN IDs to VLAN Demux Interfaces

	Single-Tag Framing	Dual-Tag Framing
Statement Format	vlan-id *vlan-id*;	vlan-tags outer *tpid*.<*vlan-id*> inner *tpidvlan-id*;

You can include all of the statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**
- **[edit interfaces *demux0* unit *logical-unit-number*]**

Associating a VLAN ID to a VLAN Demux Interface

A VLAN demux interface that you have associated to a particular VLAN ID receives and forwards incoming frames that contain a matching VLAN ID. You can associate a VLAN ID to a single-tag logical interface or to a dual-tagged (stacked) logical interface.

Associating a VLAN ID to a Single-Tag VLAN Demux Interface

To associate a VLAN ID to a single-tag VLAN demux interface, include the **vlan-id** statement at the **[edit interfaces *demux0* unit *logical-unit-number*]** hierarchy level:

```
vlan-id vlan-id;
```

To configure an interface to support single-tag logical interfaces, you must also include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

See Also

- *Configuring a VLAN Demultiplexing Interface*

Associating a VLAN ID to a Dual-Tag VLAN Demux Interface

To associate a VLAN ID to a dual-tag VLAN demux interface, include the **vlan-tags** statement at the **[edit interfaces *demux0* unit *logical-unit-number*]** hierarchy level:

vlan-tags inner <*tpid*.>*vlan-id* outer <*tpid*.>*vlan-id*;

To configure an interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

See Also

- 802.1Q VLANs Overview on page 244
- *Configuring a VLAN Demultiplexing Interface*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring VLAN and Extended VLAN Encapsulation

To configure encapsulation on an interface, enter the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
user@host# encapsulation type
```

The following list contains important notes regarding VLAN encapsulation:

- Starting with Junos OS Release 8.1, , Gigabit Ethernet IQ, Gigabit Ethernet PICs with small form-factor pluggable optics (SFPs), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use **flexible-ethernet-services**, **vlan-ccc** , or **vlan-vpls** encapsulation.
- Starting with Junos OS Release 9.5, aggregated Ethernet interfaces configured for VPLS can use **flexible-ethernet-services**, **vlan-ccc**, or **vlan-vpls**.
- Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs. For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.
- For flexible Ethernet services, Ethernet VLAN CCC and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.
- You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

- Gigabit Ethernet, 4-port Fast Ethernet, MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, 10-Gigabit Ethernet, and aggregated Ethernet interfaces with VLAN tagging enabled can use **extended-vlan-ccc** or **extended-vlan-vpls**, which allow 802.1Q tagging.
- For extended VLAN CCC and extended VLAN VPLS encapsulation, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.
- For extended VLAN CCC, the VLAN IDs on ingress and egress interfaces must be the same. For back-to-back connections, all VLAN IDs must be the same.

Release History Table

Release	Description
9.5	Starting with Junos OS Release 9.5, aggregated Ethernet interfaces configured for VPLS can use **flexible-ethernet-services**, **vlan-ccc**, or **vlan-vpls**.
8.1	Starting with Junos OS Release 8.1, , Gigabit Ethernet IQ, Gigabit Ethernet PICs with small form-factor pluggable optics (SFPs), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use **flexible-ethernet-services**, **vlan-ccc** , or **vlan-vpls** encapsulation.

Related Documentation

- *Configuring Interface Encapsulation on Physical Interfaces*
- 802.1Q VLANs Overview on page 244
- *Configuring VPLS Interface Encapsulation*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 VPN routing instance on a logical interface bound to a list of VLAN IDs.

- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 257
- Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 258
- Specifying the Interface to Handle Traffic for a CCC on page 258

Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement on a provider edge (PE) router:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
```

```
            vlan-id-list [vlan-id vlan-id–vlan-id]; # For single-tag
            vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id–vlan-id]; # For dual-tag
        }
        . . .
    }
}
```

You can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

See Also

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Specifying the Interface Over Which VPN Traffic Travels to the CE Router

To configure a Layer 2 VPN routing instance on a PE router, include the **instance-type** statement and specify the value **l2vpn**. To specify an interface connected to the router, include the **interface** statement and specify the VLAN-bundled logical interface:

```
instance-type l2vpn;
interface logical-interface-name;
```

You can include the statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

See Also

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Specifying the Interface to Handle Traffic for a CCC

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the following statements:

```
protocols {
    l2vpn {
        (control-word | no-control-word);
        encapsulation-type (ethernet | ethernet-vlan);
        site site-name {
            site-identifier identifier;
            interface logical-interface-name { # VLAN-bundled logical interface
                . . . interface-options . . .
            }
        }
    }
}
```


You can include the statements at the same hierarchy level at which you include the **instance-type l2vpn** and **interface** ***logical-interface-name*** statements:

- **[edit routing-instances** ***routing-instance-name*****]**
- **[edit logical-systems** ***logical-system-name*** **routing-instances** ***routing-instance-name*****]**

To enable a Layer 2 VPN routing instance on a PE router, include the **l2vpn** statement. For more information, see the *Junos OS VPNs Library for Routing Devices*.

The **encapsulation-type** statement specifies the Layer 2 protocol used for traffic from the customer edge (CE) router. If the Layer 2 VPN routing instance is being connected to a single-tag Layer 2 circuit, specify **ethernet** as the encapsulation type. If the Layer 2 VPN routing instance is being connected to a dual-tag Layer 2 circuit, specify **ethernet-vlan** as the encapsulation type.

To specify the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the **interface** statement and specify the VLAN-bundled logical interface.

See Also

- 802.1Q VLANs Overview on page 244
- Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 259

- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface **ge-1/0/5.0** bundles a list of VLAN IDs, and the logical interface **ge-1/1/1.0** supports IPv4 traffic using IP address 10.30.1.130 and can participate in an MPLS path.

```
[edit interfaces]
ge-1/0/5 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 { # VLAN-bundled logical interface
    vlan-id-list [513 516 520-525];
  }
}
  ge-1/1/1 {
    unit 0 {
      family inet {
        address 10.30.1.1/30;
      }
      family mpls;
    }
  }
```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance:

```
[edit protocols]
```

```
rsvp {
    interface all;
    interface lo0.0;
}
mpls {
    label-switched-path lsp {
        to 10.255.69.128;
    }
    interface all;
}
bgp {
    group g1 {
        type internal;
        local-address 10.255.69.96;
        family l2vpn {
            signaling;
        }
        neighbor 10.255.69.128;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-1/1/1.0;
    }
}
```

The following configuration shows that the VLAN-bundled logical interface is the interface over which VPN traffic travels to the CE router and handles traffic for a CCC to which the VPN connects.

```
[edit routing-instances]
red {
    instance-type l2vpn;
    interface ge-1/0/5.0; # VLAN-bundled logical interface
    route-distinguisher 10.255.69.96:100;
    vrf-target target:1:1;
    protocols {
        l2vpn {
            encapsulation-type ethernet; # For single-tag VLAN logical interface
            site CE_ultima {
                site-identifier 1;
                interface ge-1/0/5.0;
            }
        }
    }
}
```


NOTE: Because the VLAN-bundled logical interface supports single-tag frames, Ethernet is the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN.

However, with Ethernet encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Specifying the Interface Over Which VPN Traffic Travels to the CE Router

To configure a Layer 2 VPN routing instance on a PE router, include the **instance-type** statement and specify the value **l2vpn**. To specify an interface connected to the router, include the **interface** statement and specify the VLAN-bundled logical interface:

```
instance-type l2vpn;
interface logical-interface-name;
```

You can include the statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Access Mode on a Logical Interface

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified bridge domain. A logical interface configured to accept untagged packets is called an *access interface* or *access port*. Access interface configuration is supported on MX Series routers only.

To configure access mode on a logical interface, use the **interface-mode access** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family bridge]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]** hierarchy level.

When an untagged packet is received on an access interface, the packet is accepted, the configured VLAN ID is added to the packet, and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

The following example configures a logical interface as an access port with a VLAN ID of 20:

```
[edit interfaces ge-1/2/0]
unit 0 {
  family bridge {
    interface-mode access;
    vlan-id 20;
  }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Logical Interface for Trunk Mode

As an alternative to configuring a logical interface for each VLAN, enterprise network administrators can configure a single logical interface to accept untagged packets or packets tagged with any VLAN ID specified in a list of VLAN IDs. Using a VLAN ID list conserves switch resources and simplifies configuration. A logical interface configured to accept packets tagged with any VLAN ID specified in a list is called a *trunk interface* or *trunk port*. Trunk interface configuration is supported on MX Series routers only. Trunk interfaces support integrated routing and bridging (IRB).

To configure a logical interface to accept any packet tagged with a VLAN ID that matches the list of VLAN IDs, include the **interface-mode** statement and specify the **trunk** option:

```
interface-mode trunk;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family bridge]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]**

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the VLAN ID List for a Trunk Interface

To configure the list of VLAN IDs to be accepted by the trunk port, include the **vlan-id-list** statement and specify the list of VLAN IDs. You can specify individual VLAN IDs with a space separating the ID numbers, specify a range of VLAN IDs with a dash separating the ID numbers, or specify a combination of individual VLAN IDs and a range of VLAN IDs.

vlan-id-list [*number number-number*];

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]**

- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]**

When a packet is received that is tagged with a VLAN ID specified in the trunk interface list of VLAN IDs, the packet is accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.

When a packet is received that is tagged with a VLAN ID not specified in the trunk interface list of VLAN IDs, the native VLAN ID is pushed in front of the existing VLAN tag or tags and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

When an untagged packet is received on a trunk interface, the native VLAN ID is added to the packet and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

A bridge domain configured with a matching VLAN ID must be configured before the trunk interface is configured. To learn more about configuring bridge domains, see the *Junos Routing Protocols Configuration Guide*.

Related Documentation

- 802.1Q VLANs Overview on page 244

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Trunk Interface on a Bridge Network

On MX Series routers, you can configure a trunk interface on a bridge network.

The following output sample shows trunk port configuration on a bridge network:

```
user@host# run show interfaces
ge-0/0/0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
}
ge-2/0/0 {
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 1-200;
        }
    }
}
ge-2/0/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
}
```


If you want **igmp-snooping** to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge domain. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed as shown after the sample configuration below:

```
user@host# run show configuration
interfaces {
    ge-5/1/1 {
        flexible-vlan-tagging;
        native-vlan-id 1;
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list 401;
            }
        }
    }
    irb {
        unit 401 {
            family inet {
                address 192.168.2.2/27;
            }
        }
    }
}
protocols {
    igmp {
        interface all;
    }
}
bridge-domains {
    VLAN-401 {
        vlan-id 401;
        routing-interface irb.401;
        protocols {
            igmp-snooping;
        }
    }
}
```

```
user@host# commit
[edit bridge-domains]
  'VLAN-401'
    IGMP Snooping not supported with IRB and trunk mode interface ge-5/1/1.0
commit complete
```

To achieve IGMP snooping for a bridge domain, you should use such a configuration as shown in the following example:

```
user@host# run show configuration
interfaces {
    ge-0/0/1 {
        flexible-vlan-tagging;
        native-vlan-id 1;
        encapsulation flexible-ethernet-services;
        unit 0 {
            encapsulation vlan-bridge;
            vlan-id 401;
```

```
        }
    }
    irb {
        unit 401 {
            family inet {
                address 192.168.2.2/27;
            }
        }
    }
}
protocols {
    igmp {
        interface all;
    }
}
bridge-domains {
    VLAN-401 {
        vlan-id 401;
        interface ge-0/0/1.0;
        routing-interface irb.401;
        protocols {
            igmp-snooping;
        }
    }
}
```

```
user@host# commit
commit complete
```

Related Documentation

- 802.1Q VLANs Overview on page 244
- *interface-mode*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement on a provider edge (PE) router:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      vlan-id-list [vlan-id vlan-id–vlan-id]; # For single-tag
      vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id–vlan-id]; # For dual-tag
    }
    ...
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit]**

- **[edit logical-systems *logical-system-name*]**

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      encapsulation vlan–ccc; # Required for single-tag
      vlan-id-list [vlan-id vlan-id–vlan-id]; # For single-tag
      vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id–vlan-id]; # For dual-tag
    }
    . . .
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

For a single-tag logical interface, include the **encapsulation** statement and specify **vlan-ccc** so that CCC circuit encapsulation is used inside the Layer 2 circuit.

NOTE: In the case of a dual-tag logical interface, the Junos OS automatically uses the vlan-ccc encapsulation type.

Related Documentation

- 802.1Q VLANs Overview on page 244
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 268
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 circuit on a logical interface bound to a list of VLAN IDs.

- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 267
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 268

Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      encapsulation vlan–ccc; # Required for single-tag
      vlan-id-list [vlan-id vlan-id–vlan-id]; # For single-tag
      vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id–vlan-id]; # For dual-tag
    }
    . . .
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

For a single-tag logical interface, include the **encapsulation** statement and specify **vlan-ccc** so that CCC circuit encapsulation is used inside the Layer 2 circuit.

NOTE: In the case of a dual-tag logical interface, the Junos OS automatically uses the **vlan-ccc** encapsulation type.

See Also

- 802.1Q VLANs Overview on page 244
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 268

- *Ethernet Interfaces Feature Guide for Routing Devices*

Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the following statements:

```
l2circuit {
  neighbor address {
    interface logical-interface-name {
      virtual-circuit-id number;
      no-control-word;
    }
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

To enable a Layer 2 circuit, include the **l2circuit** statement.

To configure the router as a neighbor for a Layer 2 circuit, specify the neighbor address using the **neighbor** statement.

To specify the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the **interface** statement and specify the VLAN-bundled logical interface.

See Also

- 802.1Q VLANs Overview on page 244
- Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 268
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 266
- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface **ge-1/0/5.0** bundles a list of VLAN IDs, and the logical interface **ge-1/1/1.0** supports IPv4 traffic using IP address 10.30.1.1/30 and can participate in an MPLS path.

```
[edit interfaces]
ge-1/0/5 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 { # VLAN-bundled logical interface
    vlan-id-list [513 516 520-525];
  }
}
ge-1/1/1 {
  unit 0 {
    family inet {
      address 10.30.1.1/30;
```



```
        }
        family mpls;
    }
}
```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance, and shows that the VLAN-bundled logical interface handles traffic for a CCC to which the Layer 2 circuit connects:

```
[edit protocols]
rsvp {
    interface all;
    interface lo0.0;
}
mpls {
    label-switched-path lsp {
        to 10.255.69.128;
    }
    interface all;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-1/1/1.0;
    }
}
ldp {
    interface ge-1/1/1.0;
    interface ge-1/0/5.0; # VLAN-bundled logical interface
    interface lo0.0;
}
l2circuit {
    neighbor 10.255.69.128 {
        interface ge-1/0/5.0 { # VLAN-bundled logical interface
            virtual-circuit-id 3;
            no-control-word;
        }
    }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Ethernet Interfaces Feature Guide for Routing Devices*

Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs

For MX Series routers, you can bind a list of VLAN IDs to a logical interface, configure a Layer 2 VPN routing instance or Layer 2 circuit on the logical interface, and then use the logical interface to configure a circuit cross-connect (CCC) to another Layer 2 VPN routing instance or Layer 2 circuit.

A CCC allows you to configure transparent connections between two circuits so that packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. You configure a CCC by connecting circuit interfaces of the same type. For more information, see *Circuit and Translational Cross-Connects Overview*.

NOTE: **The Junos OS supports binding of Ethernet logical interfaces to lists of VLAN IDs on MX Series routers only. For all other routers, you can bind an Ethernet logical interface to only a single VLAN ID or to a single range of VLAN IDs.**

The following configuration guidelines apply to bundling lists of VLAN IDs to Ethernet logical interfaces used to configure CCCs:

Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs

To enable a physical interface to support VLAN-bundled logical interfaces that you will use to configure a CCC, you must specify one of the following physical link-layer encapsulation types as the value of the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | flexible-ethernet-services);
```

- **extended-vlan-ccc**—For Ethernet interfaces with standard TPID tagging.
- **flexible-ethernet-services**—For supported Gigabit Ethernet interfaces for which you want to configure multiple per-unit Ethernet encapsulations.

For more information about configuring the encapsulation on a physical interface, see *Configuring Interface Encapsulation on Physical Interfaces*.

Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs

For VLAN-bundled logical interfaces that you use to configure a CCC, specific logical link-layer encapsulation types are used inside the circuits themselves.

Table 17 on page 271 describes the logical link-layer encapsulation types used within circuits connected using VLAN-bundled logical interfaces of the same type.

Table 17: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces

Encapsulation Inside the Circuit	Layer 2 Circuit Joined by Configuring an Interface-to-Interface CCC Connection	
	Layer 2 VPN Routing Instance	Layer 2 Circuit
Syntax	encapsulation-type (ethernet | ethernet-vlan);	encapsulation vlan-ccc;
Hierarchy Level	[edit routing-instances *routing-instance-name* protocols l2vpn], [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]	[edit interfaces *ethernet-interface-name* unit *logical-unit-number*], [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]
Usage Guidelines	See the *Junos OS VPNs Library for Routing Devices*.	See *Configuring Interface Encapsulation on Logical Interfaces*, *Circuit and Translational Cross-Connects Overview*, and *Defining the Encapsulation for Switching Cross-Connects*.
For a Single-Tag Logical Interface	The MX Series router automatically uses **ethernet** as the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN. **NOTE:** With **ethernet** encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.	Configure the MX Series router to use **vlan-ccc** as the logical link-layer encapsulation type.
For a Dual-Tag Logical Interface	Configure the MX Series router to use **ethernet-vlan** as the Layer 2 protocol to encapsulate incoming traffic. With **ethernet-vlan** encapsulation, circuit signal processing checks that the VLAN ID list is the same at both ends of the CCC connection. If a VLAN ID list mismatch is detected, you can view the error condition in the **show interfaces** command output.	The MX Series router automatically uses **vlan-ccc** as the logical link-layer encapsulation type, regardless of the value configured.

Related Documentation

- 802.1Q VLANs Overview on page 244
- Binding VLAN IDs to Logical Interfaces on page 251

- *Defining the Encapsulation for Switching Cross-Connects*

Specifying the Interface to Handle Traffic for a CCC

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the following statements:

```
protocols {
  l2vpn {
    (control-word | no-control-word);
    encapsulation-type (ethernet | ethernet-vlan);
    site site-name {
```

```
        site-identifier identifier;
        interface logical-interface-name { # VLAN-bundled logical interface
          . . . interface-options . . .
        }
      }
    }
  }
```

You can include the statements at the same hierarchy level at which you include the **instance-type l2vpn** and **interface *logical-interface-name*** statements:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

To enable a Layer 2 VPN routing instance on a PE router, include the **l2vpn** statement. For more information, see the *Junos OS VPNs Library for Routing Devices*.

The **encapsulation-type** statement specifies the Layer 2 protocol used for traffic from the customer edge (CE) router. If the Layer 2 VPN routing instance is being connected to a single-tag Layer 2 circuit, specify **ethernet** as the encapsulation type. If the Layer 2 VPN routing instance is being connected to a dual-tag Layer 2 circuit, specify **ethernet-vlan** as the encapsulation type.

To specify the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the **interface** statement and specify the VLAN-bundled logical interface.

Related Documentation

- 802.1Q VLANs Overview on page 244
- Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 259
- *Ethernet Interfaces Feature Guide for Routing Devices*

Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the following statements:

```
l2circuit {
  neighbor address {
    interface logical-interface-name {
      virtual-circuit-id number;
      no-control-word;
    }
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

To enable a Layer 2 circuit, include the **l2circuit** statement.

To configure the router as a neighbor for a Layer 2 circuit, specify the neighbor address using the **neighbor** statement.

To specify the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the **interface** statement and specify the VLAN-bundled logical interface.

Related Documentation

- 802.1Q VLANs Overview on page 244
- Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 268
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 266
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 12

Configuring Private VLANs

Understanding Private VLANs

VLANs limit broadcasts to specified users. Private VLANs (PVLANs) take this concept a step further by limiting communication within a VLAN. PVLANs accomplish this by restricting traffic flows through their member switch ports (which are called *private ports*) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. The uplink trunk port or link aggregation group (LAG) is usually connected to a router, firewall, server, or provider network. Each PVLAN typically contains many private ports that communicate only with a single uplink port, thereby preventing the ports from communicating with each other.

PVLANs provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs (community VLANs and an isolated VLAN) inside a primary VLAN. Ports within the same community VLAN can communicate with each other. Ports within an isolated VLAN can communicate *only* with a single uplink port.

Just like regular VLANs, PVLANs are isolated on Layer 2 and require one of the following options to route Layer 3 traffic among the secondary VLANs:

- A promiscuous port connection with a router
- A routed VLAN interface (RVI)

NOTE: To route Layer 3 traffic among secondary VLANs, a PVLAN needs only one of the options mentioned above. If you use an RVI, you can still implement a promiscuous port connection to a router with the promiscuous port set up to handle only traffic that enters and exits the PVLAN.

PVLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts. Service providers use PVLANs to keep their customers isolated from each other. Another typical use for a PVLAN is to provide per-room Internet access in a hotel.

NOTE: You can configure a PVLAN to span switches that support PVLANs.

This topic explains the following concepts regarding PVLANs on EX Series switches:

Why Use PVLANs

PVLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between hosts.The need to segregate a single VLAN is particularly useful in the following deployment scenarios:

- Server farms—A typical Internet service provider uses a server farm to provide Web hosting for numerous customers. Locating the various servers within a single server farm provides ease of management. Security concerns arise if all servers are in the same VLAN because Layer 2 broadcasts go to all servers in the VLAN.
- Metropolitan Ethernet networks—A metro service provider offers Layer 2 Ethernet access to assorted homes, rental communities, and businesses. The traditional solution of deploying one VLAN per customer is not scalable and is difficult to manage, leading

to potential waste of IP addresses. PVLANs provide a more secure and more efficient solution.

Typical Structure and Primary Application of PVLANs

A PVLAN can be configured on a single switch or can be configured to span multiple switches. The types of domains and ports are:

- Primary VLAN—The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID) for the complete PVLAN. The primary PVLAN can contain multiple secondary VLANs (one isolated VLAN and multiple community VLANs).
- Isolated VLAN/isolated port—A primary VLAN can contain only one isolated VLAN. An interface within an isolated VLAN can forward packets only to a promiscuous port or the Inter-Switch Link (ISL) port. An isolated interface cannot forward packets to another isolated interface; and an isolated interface cannot receive packets from another isolated interface. If a customer device needs to have access *only* to a gateway router, the device must be attached to an isolated trunk port.
- Community VLAN/community port—You can configure multiple community VLANs within a single PVLAN. An interface within a specific community VLAN can establish Layer 2 communications with any other interface that belongs to the same community VLAN. An interface within a community VLAN can also communicate with a promiscuous port or the ISL port. If you have, for example, two customer devices that you need to isolate from other customer devices but that must be able to communicate with one another, use community ports.
- Promiscuous port—A promiscuous port has Layer 2 communications with all interfaces in the PVLAN, regardless of whether an interface belongs to an isolated VLAN or a community VLAN. A promiscuous port is a member of the primary VLAN but is not included within any secondary subdomain. Layer 3 gateways, DHCP servers, and other trusted devices that need to communicate with endpoint devices are typically connected to a promiscuous port.
- Inter-Switch Link (ISL)—An ISL is a trunk port that connects multiple switches in a PVLAN and contains two or more VLANs. It is required only when a PVLAN spans multiple switches.

The configured PVLAN is the *primary* domain (primary VLAN). Within the PVLAN, you configure *secondary* VLANs, which become subdomains nested within the primary domain. A PVLAN can be configured on a single switch or can be configured to span multiple switches. The PVLAN shown in Figure 18 on page 278 includes two switches, with a primary PVLAN domain and various subdomains.

Figure 18: Subdomains in a PVLAN

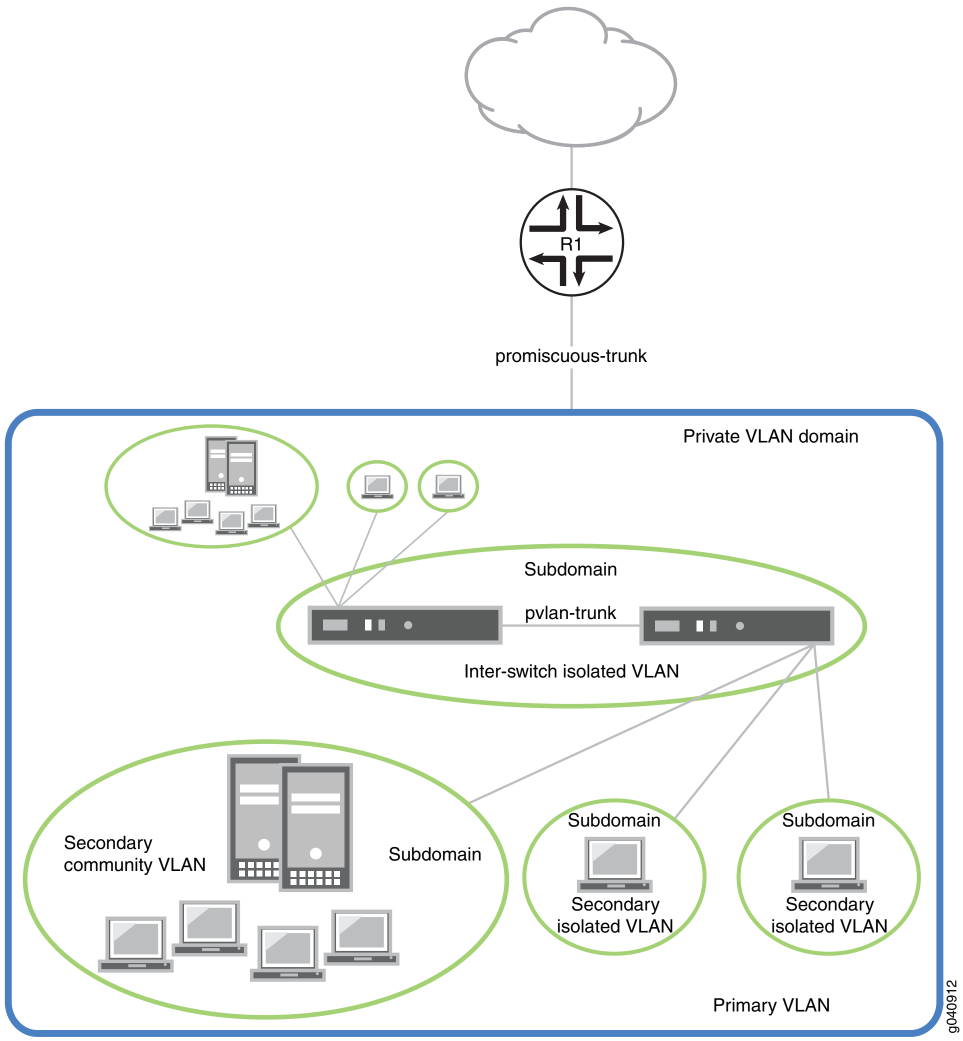

As shown in Figure 20 on page 280, a PVLAN has only one primary domain and multiple secondary domains. The types of domains are:

- Primary VLAN—VLAN used to forward frames downstream to isolated and community VLANs. The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID) for the complete PVLAN. The primary PVLAN can contain multiple secondary VLANs (one isolated VLAN and multiple community VLANs).
- Secondary isolated VLAN—VLAN that receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN. The isolated VLAN is a secondary VLAN nested within the primary VLAN. A primary VLAN can contain only one isolated VLAN. An interface within an isolated VLAN (isolated interface) can forward packets only to a promiscuous port or the PVLAN trunk port. An isolated interface cannot forward packets to another isolated interface; nor can an isolated interface receive packets from another isolated interface. If a customer device needs to have access *only* to a router, the device must be attached to an isolated trunk port.

- Secondary interswitch isolated VLAN—VLAN used to forward isolated VLAN traffic from one switch to another through PVLAN trunk ports. 802.1Q tags are required for interswitch isolated VLANs because IEEE 802.1Q uses an internal tagging mechanism by which a trunking device inserts a 4-byte VLAN frame identification tab into the packet header. An interswitch isolated VLAN is a secondary VLAN nested within the primary VLAN.
- Secondary community VLAN—VLAN used to transport frames among members of a community (a subset of users within the VLAN) and to forward frames upstream to the primary VLAN. A community VLAN is a secondary VLAN nested within the primary VLAN. You can configure multiple community VLANs within a single PVLAN. An interface within a specific community VLAN can establish Layer 2 communications with any other interface that belongs to the same community VLAN. An interface within a community VLAN can also communicate with a promiscuous port or the PVLAN trunk port.

Figure 19 on page 279 shows a PVLAN spanning multiple switches, where the primary VLAN (**100**) contains two community domains (**300** and **400**) and one interswitch isolated domain.

Figure 19: PVLAN Spanning Multiple Switches

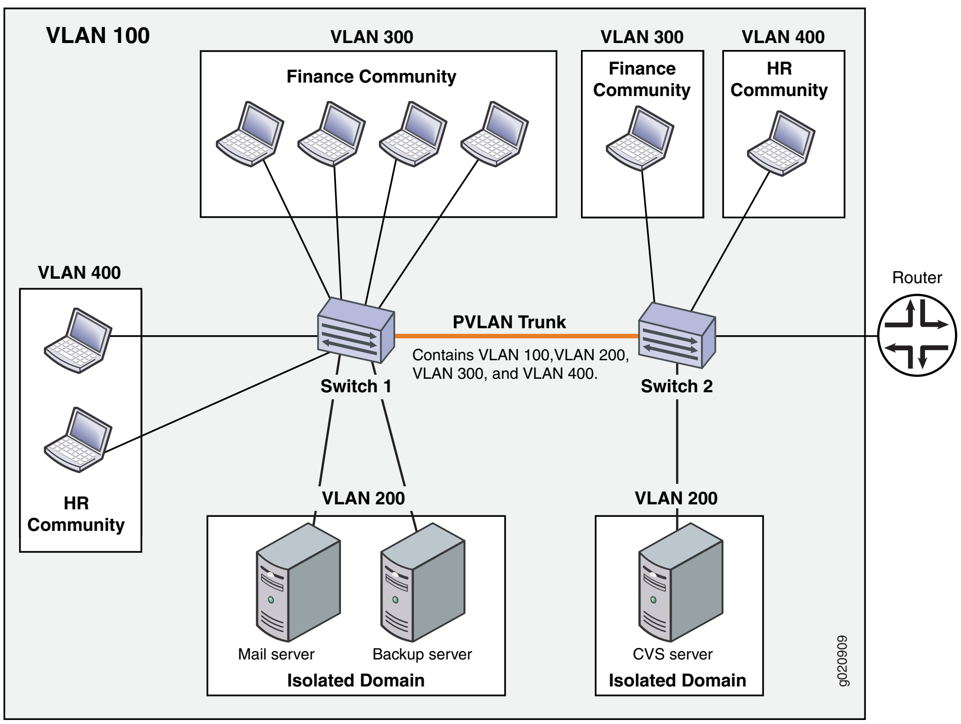

NOTE: Primary and secondary VLANs count against the limit of 4089 VLANs supported on the QFX Series. For example, each VLAN in Figure 19 on page 279 counts against this limit.

Typical Structure and Primary Application of PVLANs on MX Series Routers

The configured PVLAN becomes the primary domain, and secondary VLANs become subdomains that are nested inside the primary domain. A PVLAN can be created on a single router. The PVLAN shown in Figure 20 on page 280 includes one router, with one primary PVLAN domain and multiple secondary subdomains.

Figure 20: Subdomains in a PVLAN With One Router

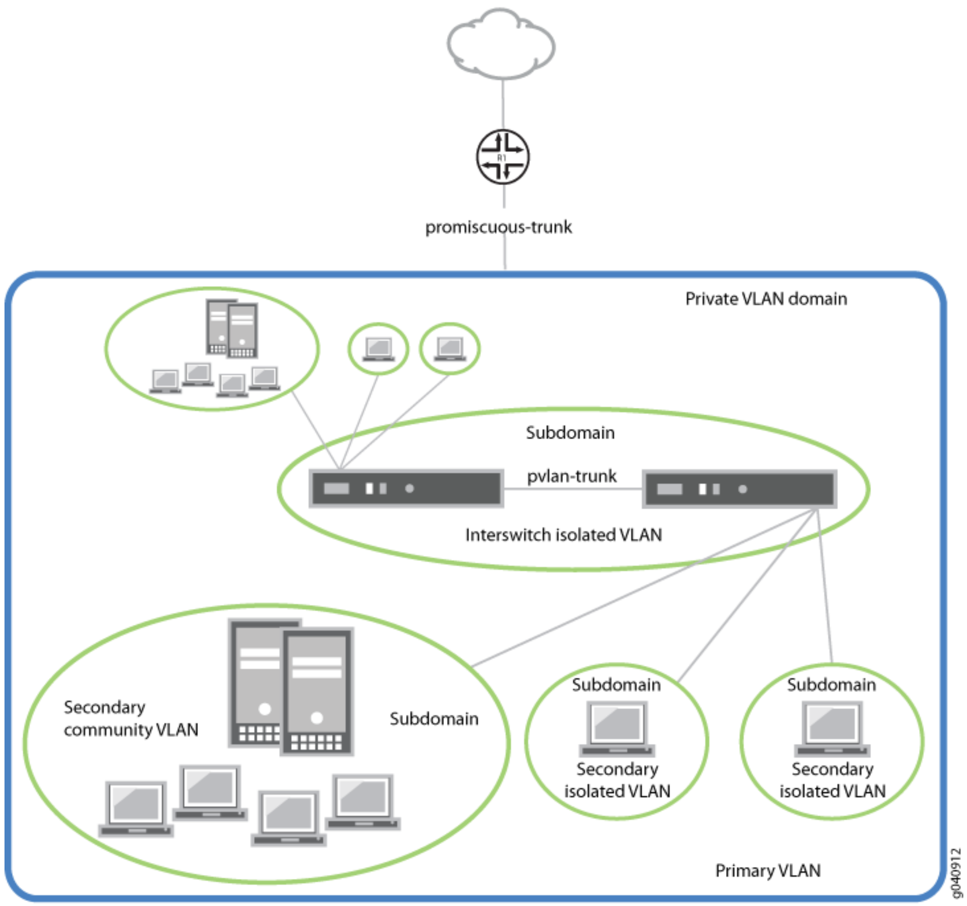

The types of domains are:

- Primary VLAN—VLAN used to forward frames downstream to isolated and community VLANs.
- Secondary isolated VLAN—VLAN that receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.
- Secondary interswitch isolated VLAN—VLAN used to forward isolated VLAN traffic from one router to another through PVLAN trunk ports.
- Secondary community VLAN—VLAN used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN.

NOTE: PVLANs are supported on MX80 routers, on MX240, MX480, and MX960 routers with DPCs in enhanced LAN mode, on MX Series routers with MPC1, MPC2, and Adaptive Services PICs.

Typical Structure and Primary Application of PVLANs on EX Series Switches

NOTE: The primary VLAN of the PVLAN is defined with an 802.1Q tag (VLAN ID) for the complete PVLAN. On EX9200 switches, each secondary VLAN must also be defined with its own separate VLAN ID.

Figure 21 on page 281 shows a PVLAN on a single switch, where the primary VLAN (VLAN **100**) contains two community VLANs (VLAN **300** and VLAN **400**) and one isolated VLAN (VLAN **50**).

Figure 21: Private VLAN on a Single EX Switch

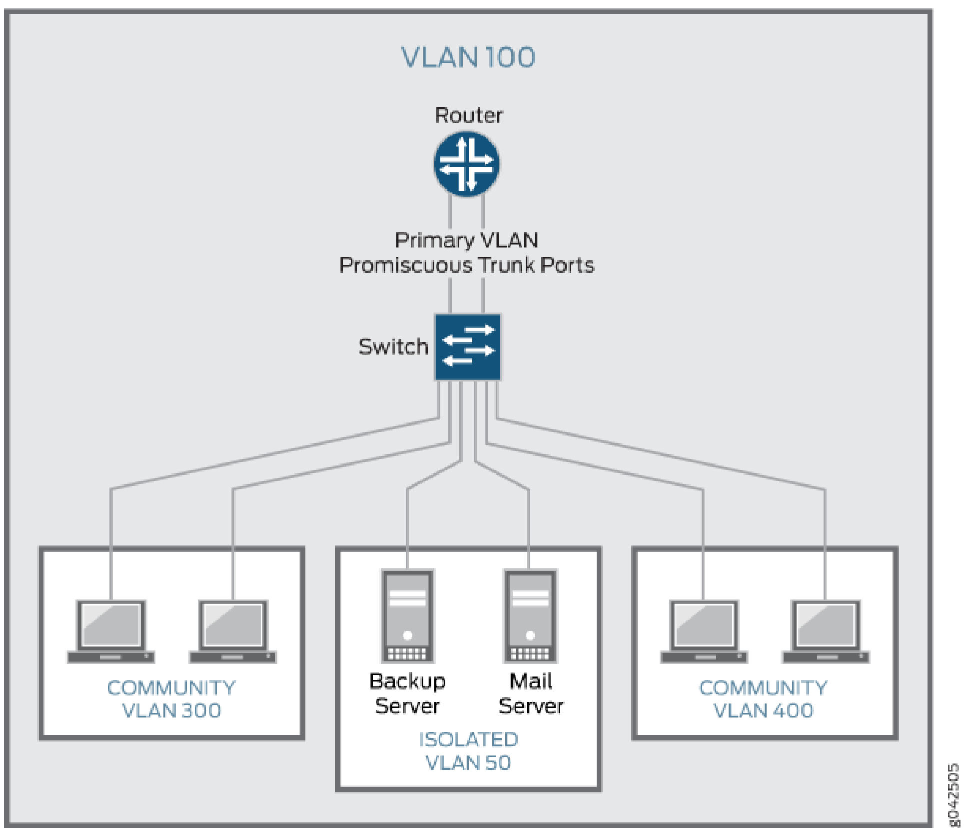

Figure 22 on page 282 shows a PVLAN spanning multiple switches, where the primary VLAN (VLAN **100**) contains two community VLANs (VLAN **300** and VLAN **400**) and one isolated VLAN (VLAN 200). It also shows that Switches 1 and 2 are connected through an interswitch link (PVLAN trunk link).

Figure 22: PVLAN Spanning Multiple EX Series Switches

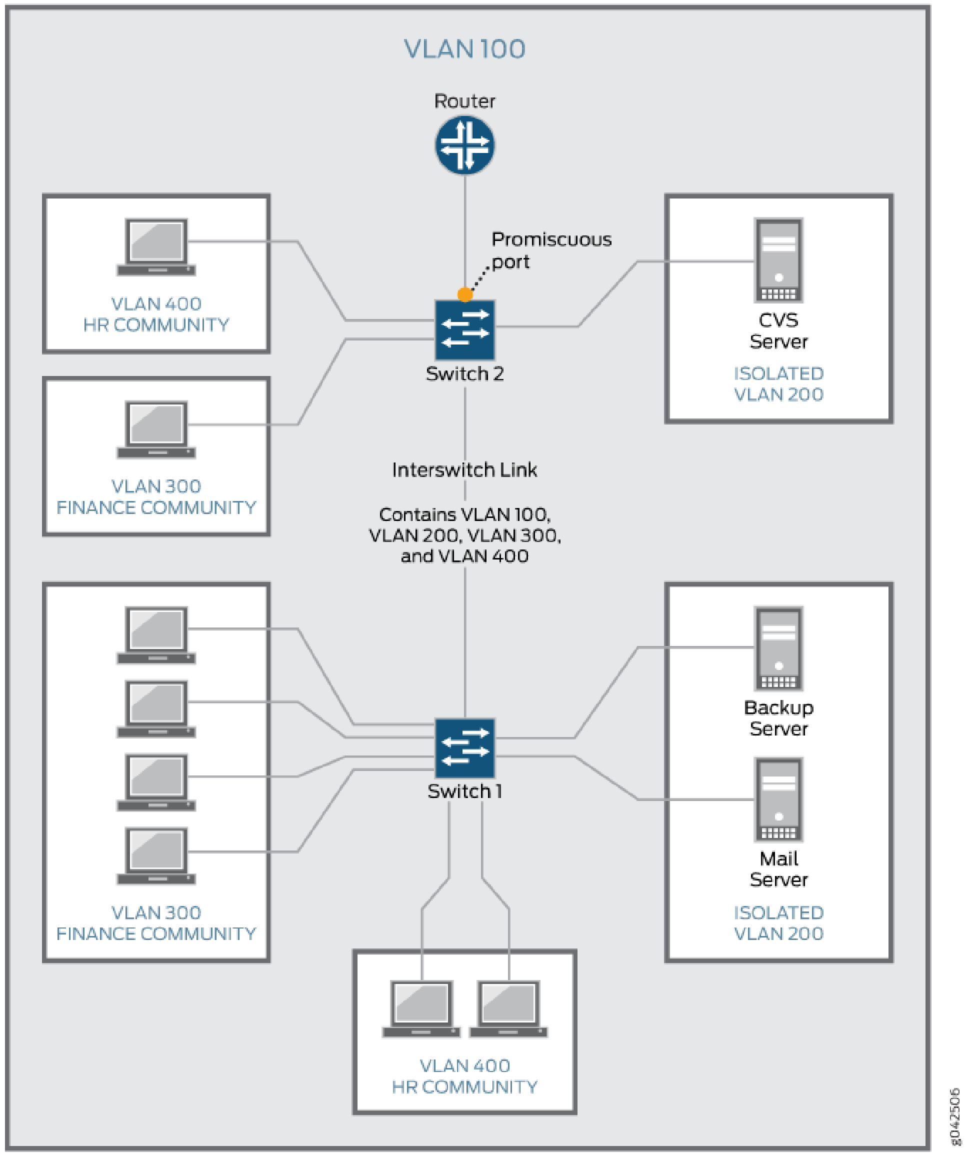

Also, the PVLANs shown in Figure 21 on page 281 and Figure 22 on page 282 use a promiscuous port connected to a router as the means to route Layer 3 traffic among the community and isolated VLANs. Instead of using the promiscuous port connected to a router, you can configure an RVI on the switch in Figure 21 on page 281 or one of the switches shown in Figure 22 on page 282 (on some EX switches).

To route Layer 3 traffic between isolated and community VLANs, you must either connect a router to a promiscuous port, as shown in Figure 21 on page 281 and Figure 22 on page 282, or configure an RVI.

If you choose the RVI option, you must configure one RVI for the primary VLAN in the PVLAN domain. This RVI serves the entire PVLAN domain regardless of whether the domain includes one or more switches. After you configure the RVI, Layer 3 packets received by the secondary VLAN interfaces are mapped to and routed by the RVI.

When setting up the RVI, you must also enable proxy Address Resolution Protocol (ARP) so that the RVI can handle ARP requests received by the secondary VLAN interfaces.

For information about configuring PVLANs on a single switch and on multiple switches, see *Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)* and *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*, respectively. For information about configuring an RVI, see *Configuring a Routed VLAN Interface in a Private VLAN on an EX Series Switch (CLI Procedure)*.

Routing Between Isolated and Community VLANs

To route Layer 3 traffic between isolated and community VLANs, you must connect an external router or switch to a trunk port of the primary VLAN. The trunk port of the primary VLAN is a *promiscuous* port; therefore, it can communicate with *all* the ports in the PVLAN.

PVLANs Use 802.1Q Tags to Identify Packets

When packets are marked with a customer-specific 802.1Q tag, that tag identifies ownership of the packets for any switch or router in the network. Sometimes, 802.1Q tags are needed within PVLANs to keep track of packets from different subdomains. Table 18 on page 283 indicates when a VLAN 802.1Q tag is needed on the primary VLAN or on secondary VLANs.

Table 18: When VLANs in a PVLAN Need 802.1Q Tags

	On a Single Switch	On Multiple Switches
Primary VLAN	Specify an 802.1Q tag by setting a VLAN ID.	Specify an 802.1Q tag by setting a VLAN ID.
Secondary VLAN	No tag needed on VLANs.	VLANs need 802.1Q tags: • Specify an 802.1Q tag for each community VLAN by setting a VLAN ID. • Specify the 802.1Q tag for an isolation VLAN ID by setting an isolation ID.

PVLANs Use IP Addresses Efficiently

PVLANs provide IP address conservation and efficient allocation of IP addresses. In a typical network, VLANs usually correspond to a single IP subnet. In PVLANs, the hosts in all secondary VLANs belong to the same IP subnet because the subnet is allocated to the primary VLAN. Hosts within the secondary VLAN are assigned IP addresses based on IP subnets associated with the primary VLAN, and their IP subnet masking information reflects that of the primary VLAN subnet. However, each secondary VLAN is a separate broadcast domain.

PVLAN Port Types and Forwarding Rules

PVLANs can use up to six different port types. The network depicted in Figure 19 on page 279 uses a promiscuous port to transport information to the router, community ports to connect the finance and HR communities to their respective switches, isolated ports to connect the servers, and a PVLAN trunk port to connect the two switches. PVLAN ports have different restrictions:

- Promiscuous trunk port—A promiscuous port has Layer 2 communications with all the interfaces that are in the PVLAN, regardless of whether the interface belongs to an isolated VLAN or a community VLAN. A promiscuous port is a member of the primary VLAN, but is not included within one of the secondary subdomains. Layer 3 gateways, DHCP servers, and other trusted devices that need to communicate with endpoint devices are typically connected to a promiscuous port.
- PVLAN trunk link—The PVLAN trunk link, which is also known as the interswitch link, is required only when a PVLAN is configured to span multiple switches. The PVLAN trunk link connects the multiple switches that compose the PVLAN.
- PVLAN trunk port—A PVLAN trunk port is required in multiswitch PVLAN configurations to span the switches. The PVLAN trunk port is a member of all VLANs within the PVLAN (that is, the primary VLAN, the community VLANs, and the interswitch isolated VLAN), and it carries traffic from the primary VLAN and all secondary VLANs. It can communicate with all ports other than the isolated ports.

  Communication between a PVLAN trunk port and an isolated port is usually unidirectional. A PVLAN trunk port's membership in the interswitch isolated VLAN is egress-only, meaning that an isolated port can forward packets to a PVLAN trunk port, but a PVLAN trunk port does not forward packets to an isolated port (unless the packets ingressed on a promiscuous access port and are therefore being forwarded to all the secondary VLANs in the same primary VLAN as the promiscuous port).
- Secondary VLAN trunk port (not shown)—Secondary trunk ports carry secondary VLAN traffic. For a given private VLAN, a secondary VLAN trunk port can carry traffic for only one secondary VLAN. However, a secondary VLAN trunk port can carry traffic for multiple secondary VLANs as long as each secondary VLAN is a member of a different primary VLAN. For example, a secondary VLAN trunk port can carry traffic for a community VLAN that is part of primary VLAN pvlan100 and also carry traffic for an isolated VLAN that is part of primary VLAN pvlan400.
- Community port—Community ports communicate among themselves and with their promiscuous ports. Community ports serve only a select group of users. These interfaces are separated at Layer 2 from all other interfaces in other communities or isolated ports within their PVLAN.
- Isolated access port—Isolated ports have Layer 2 connectivity only with promiscuous ports and PVLAN trunk ports—an isolated port cannot communicate with another isolated port even if these two ports are members of the same isolated VLAN (or interswitch isolated VLAN) domain. Typically, a server, such as a mail server or a backup server, is connected on an isolated port. In a hotel, each room would typically be

connected on an isolated port, meaning that room-to-room communication is not possible, but each room can access the Internet on the promiscuous port.

- Promiscuous access port (not shown)—These ports carry untagged traffic. Traffic that ingresses on a promiscuous access port is forwarded to all secondary VLAN ports on the device. If traffic ingresses into the device on a VLAN-enabled port and egresses on a promiscuous access port, the traffic is untagged on egress. If tagged traffic ingresses on a promiscuous access port, the traffic is discarded.
- Interswitch link port—An interswitch link (ISL) port is a trunk port that connects two routers when a PVLAN spans those routers. The ISL port is a member of all VLANs within the PVLAN (that is, the primary VLAN, the community VLANs, and the isolated VLAN).

  Communication between an ISL port and an isolated port is unidirectional. An ISL port's membership in the interswitch isolated VLAN is egress-only, meaning that incoming traffic on the ISL port is never assigned to the isolated VLAN. An isolated port can forward packets to a PVLAN trunk port, but a PVLAN trunk port cannot forward packets to an isolated port. Table 20 on page 285 summarizes whether Layer 2 connectivity exists between the different types of ports.

Table 19 on page 285 summarizes Layer 2 connectivity between the different types of ports within a PVLAN on EX Series switches that support ELS.

Table 19: PVLAN Ports and Layer 2 Forwarding on EX Series switches that support ELS

From Port Type	To Isolated Ports?	To Promiscuous Ports?	To Community Ports?	To Inter-Switch Link Port?
Isolated	Deny	Permit	Deny	Permit
Promiscuous	Permit	Permit	Permit	Permit
Community 1	Deny	Permit	Permit	Permit

Table 20: PVLAN Ports and Layer 2 Connectivity

Port Type	Promiscuous Trunk	PVLAN Trunk	Secondary Trunk	Community	Isolated Access	Promiscuous access
Promiscuous trunk	Yes	Yes	Yes	Yes	Yes	Yes
PVLAN trunk	Yes	Yes	Yes	Yes—same community only	Yes	Yes
Secondary Trunk	Yes	Yes	No	Yes	No	Yes
Community	Yes	Yes	Yes	Yes—same community only	No	Yes
Isolated access	Yes	Yes—unidirectional only	No	No	No	Yes

Table 20: PVLAN Ports and Layer 2 Connectivity (continued)

Port Type	Promiscuous Trunk	PVLAN Trunk	Secondary Trunk	Community	Isolated Access	Promiscuous access
Promiscuous access	Yes	Yes	Yes	Yes	Yes	No

Table 21 on page 286 summarizes whether or not Layer 2 connectivity exists between the different types of ports within a PVLAN.

Table 21: PVLAN Ports and Layer 2 Connectivity on EX Series Switches without ELS Support

Port Type To: → From:↓	Promiscuous	Community	Isolated	PVLAN Trunk	RVI
Promiscuous	Yes	Yes	Yes	Yes	Yes
Community	Yes	Yes—same community only	No	Yes	Yes
Isolated	Yes	No	No	Yes NOTE: This communication is unidirectional.	Yes
PVLAN trunk	Yes	Yes—same community only	Yes NOTE: This communication is unidirectional.	Yes	Yes
RVI	Yes	Yes	Yes	Yes	Yes

As noted in Table 21 on page 286, Layer 2 communication between an isolated port and a PVLAN trunk port is unidirectional. That is, an isolated port can only send packets to a PVLAN trunk port, and a PVLAN trunk port can only receive packets from an isolated port. Conversely, a PVLAN trunk port cannot send packets to an isolated port, and an isolated port cannot receive packets from a PVLAN trunk port.

NOTE: If you enable **no-mac-learning** on a primary VLAN, all isolated VLANs (or the interswitch isolated VLAN) in the PVLAN inherit that setting. However, if you want to disable MAC address learning on any community VLANs, you must configure **no-mac-learning** on each of those VLANs.

Creating a PVLAN

The flowchart shown in Figure 23 on page 287 gives you a general idea of the process for creating PVLANs. If you complete your configuration steps in the order shown, you will not violate these PVLAN rules. (In the PVLAN rules, configuring the PVLAN trunk port applies only to a PVLAN that spans multiple routers.)

- The primary VLAN must be a tagged VLAN.
- If you are going to configure a community VLAN ID, you must first configure the primary VLAN.
- If you are going to configure an isolation VLAN ID, you must first configure the primary VLAN.

NOTE: **Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.**

Configuring a VLAN on a single router is relatively simple, as shown in Figure 23 on page 287.

Figure 23: Configuring a PVLAN on a Single Switch

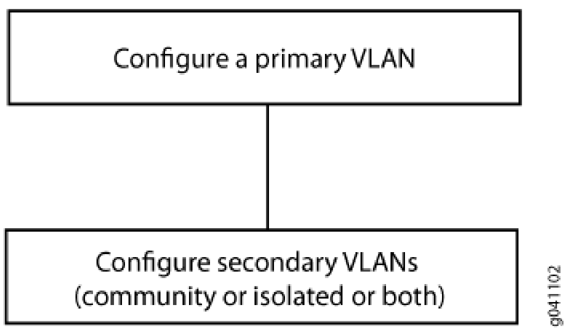

Configuring a primary VLAN consists of these steps:

1. Configure the primary VLAN name and 802.1Q tag.
2. Set **no-local-switching** on the primary VLAN.
3. Configure the promiscuous trunk port and access ports.
4. Make the promiscuous trunk and access ports members of the primary VLAN.

Within a primary VLAN, you can configure secondary community VLANs or secondary isolated VLANs or both. Configuring a secondary community VLAN consists of these steps:

1. Configure a VLAN using the usual process.

2. Configure access interfaces for the VLAN.

3. Assign a primary VLAN to the community VLAN,

Isolated VLANs are created internally when the isolated VLAN has access interfaces as members and the option **no-local-switching** is enabled on the primary VLAN.

802.1Q tags are required for interswitch isolated VLANs because IEEE 802.1Q uses an internal tagging mechanism by which a trunking device inserts a 4-byte VLAN frame identification tab into the packet header.

Trunk ports are only needed for multirouter PVLAN configurations—the trunk port carries traffic from the primary VLAN and all secondary VLANs.

Limitations of Private VLANs

The following constraints apply to private VLAN configurations:

- An access interface can belong to only one PVLAN domain, that is, it cannot participate in two different primary VLANs.
- A trunk interface can be a member of two secondary VLANs as long as the secondary VLANs are in two *different* primary VLANs. A trunk interface cannot be a member of two secondary VLANs that are in the *same* primary VLAN.
- A single region of Multiple Spanning Tree Protocol (MSTP) must be configured on all VLANs that are included within the PVLAN.
- VLAN Spanning Tree Protocol (VSTP) is not supported.
- IGMP snooping is not supported with private VLANs.
- Routed VLAN interfaces are not supported on private VLANs
- Routing between secondary VLANs in the same primary VLAN is not supported.
- Some configuration statements cannot be specified on a secondary VLAN. You can configure the following statements at the **[edit vlans *vlan-name* switch-options]** hierarchy level *only* on the primary PVLAN.
- If you want to change a primary VLAN to be a secondary VLAN, you must first change it to a normal VLAN and commit the change. For example, you would follow this procedure:
  1. Change the primary VLAN to be a normal VLAN.
  2. Commit the configuration.
  3. Change the normal VLAN to be a secondary VLAN.
  4. Commit the configuration.

Follow the same sequence of commits if you want to change a secondary VLAN to be a primary VLAN. That is, make the secondary VLAN a normal VLAN and commit that change and then change the normal VLAN to be a primary VLAN.

The following features are *not* supported on PVLANs on Junos switches with support for the ELS configuration style:

- DHCP security features (DHCP snooping, dynamic ARP inspection, IP source guard)
- Egress VLAN firewall filters
- Ethernet ring protection (ERP)
- Flexible VLAN tagging
- *global-mac-statistics*
- Integrated routing and bridging (IRB) interface
- Multicast snooping or IGMP snooping
- Multichassis link aggregation groups (MC-LAGs)
- Port mirroring
- Q-in-Q tunneling
- VLAN Spanning Tree Protocol (VSTP)
- Voice over IP (VoIP)

You can configure the following statements at the **[edit vlans *vlan-name* switch-options]** hierarchy level only on the primary PVLAN:

- *mac-table-size*
- *no-mac-learning*
- *mac-statistics*
- *interface-mac-limit*

Related Documentation

- *Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANs*
- *Creating a Private VLAN on a Single QFX Switch*

  Example: Configuring a Private VLAN on a Single Switch with ELS Support
- *Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)*
- *Creating a Private VLAN Spanning Multiple QFX Series Switches*
- *Example: Configuring a Private VLAN on a Single EX Series Switch*
- *Understanding Bridging and VLANs on Switches*
- *Example: Configuring a Private VLAN Spanning Multiple EX Series Switches*
- *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*
- Bridge Domains Setup in PVLANs on MX Series Routers on page 290

- Bridging Functions With PVLANs on page 292

Bridge Domains Setup in PVLANs on MX Series Routers

Bridge domain capabilities are used to support PVLANs on MX Series routers. Although this functionality is similar to the PVLAN mechanism on EX Series switches, the difference is that only one isolation VLAN can be configured for all isolated ports on MX routers instead of one isolation VLAN permissible per isolated port on EX Series switches.

Assume a sample deployment in which a primary VLAN named VP contains ports p1, p2, t1, t2, i1, i2, cx1, and cx2. The port types of these configured ports are as follows:

- Promiscuous ports = p1, p2
- ISL ports = t1, t2
- Isolated ports = i1, i2
- Community VLAN = Cx
- Community ports = cx1, cx2

Bridge domains are provisioned for each of the VLANs, namely, Vp, Vi, and Vcx. Assume the bridge domains to be configured as follows:

Vp—BD_primary_Vp (ports contained are p1, t1, i1, i2, cx1, cx2)

Vi—BD_isolate_Vi (ports contained are p1, t1, *i1, *i2)

Vcx—BD_community_Vcx (ports contained are p1, t1, cx1, cx2)

The bridge domains for community, primary, and isolated VLANs are automatically created by the system internally when you configure a bridge domain with a trunk interface, access interface, or interswitch link. The bridge domains contain the same VLAN ID corresponding to the VLANs. To use bridge domains for PVLANs, you must configure the following additional attributes:

- **community-vlans** option—This option is specified on all community vlans and for community BDs created internally.
- **isolated-vlan** option—This option denotes the vlan tag to be used for isolation BD created internally for each PVLAN/BD. This setting is required.
- **inter-switch-link** option with the **interface-mode trunk** statement at the **[edit interfaces interface-name family bridge]** or the **[edit interfaces interface-name unit logical-unit-number family bridge]** hierarchy level—This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member.

You can use the **vlan-id** configuration statement for PVLAN ports to identify the port role. All the logical interfaces involved in PVLANs must be configured with a VLAN ID and the Layer 2 process uses this VLAN tag to classify a port role as promiscuous, isolated, or community port by comparing this value with the VLANs configured in the PVLAN bridge domain (using the **bridge-domains** statement at the **[edit]** hierarchy level). The ISL port role is identified by the **inter-switch-link** option. The VLAN ID for ISL port is required and must be set to to the primary VLAN ID. The ISL must be a trunk interface. A list of VLAN IDs is not needed because the Layer 2 process creates such a list internally based on PVLAN bridge domain configuration. For untagged promiscuous, isolated or community, logical interfaces or ports, access mode must be used as the interface mode. For tagged promiscuous, isolated, or community interfaces, trunk mode must be specified as the interface mode.

The bridge domain interface families are enhanced to include ingress-only and egress-only association. The association for the interface family bridge domain (IFBD) is created in the following manner:

- For BD_primary_Vp, IFBD for i1, i2, cx1 and cx2 are egress only.
- BD_isolate_Vi, IFBD for p1 will be egress only and for i1 and i2 are ingress only.
- BD_community_Vcx , IFBD for p1 are egress only. VLAN translation rules ensure the following VLAN mappings to work properly:
  - VLAN mapping on promiscuous ports: On promiscuous ports, the Vlan Vi is mapped to Vlan Vp on egress interfaces. Similarly on promiscuous ports, Vcx is also be mapped to Vp.
  - VLAN mapping on isolation ports: On tagged isolated ports, the VLAN tag, Vp, is mapped to Vi on egress.
  - VLAN mapping on community ports: On tagged community ports, the VLAN tag, Vp, is mapped to Vcx on egress.

A management bridge domain for PVLAN that exists only in the Layer 2 address learning process called PBD to denote bridge domain for VLAN is used by the system. This bridge domain has the same name as the user-configured name. Under this bridge domain, one primary PVLAN bridge domain for the primary vlan, one isolation bridge domain for the isolation vlan, and one community bridge domain for each community vlan are programmed internally. You might find separate bridge domains for the PVLAN ports to be useful if you want to configure a policy for a specific community VLAN or isolation VLAN.

The management bridge domain maintains a list to include all internal bridge domains that belong to this PVLAN bridge domain. Isolation and community bridge domains contain a pointer or a flag to indicate that this bridge domain is for PVLANs and maintain the information about the primary bridge domain index and primary VLAN. All this information is available across the bridge domain interfaces that are mapped to this bridge domain. MAC learning occurs only in the primary bridge domain and the MAC forwarding entry is programmed into the primary bridge domain only. As a result, the isolation bridge domain and all community bridge domains share the same forwarding table as the primary bridge domain.

For the isolation bridge domain, BD_isolate_Vi, isolation port i1 and i2 function as a non-local-switch access port and the flood group for this bridge domain contains only the promiscuous port, p1, and ISL ports, t1 and t2.

Bridging Functions With PVLANs

This topic describes how bridging is implemented on MX Series routers that will help with understanding the unique enhancements involved in implementing PVLAN bridging procedures. Consider two ports in a bridging domain with the respective ports on different FPCs and different Packet Forwarding Engines. When a packet enters a port, the following is the flow, assuming it is a tagged packet:

1. As the starting process, a VLAN lookup is performed to determine which bridging domain the packet forms. The result of the lookup identifies the bridging domain id (bd_id), mesh group id (mg_id). With these parameters, other related information configured for this bridging domain is discovered.
2. A source MAC address (SMAC) lookup is performed to find out whether this MAC addresses is learned or not. If it is not a learned address, an MLP packet (route for flooding traffic to MAC learning chips) is sent to all the other Packet Forwarding Engines that are mapped with this bridging domain. In addition, an MLP packet is also sent to the host.
3. A destination MAC address (DMAC) lookup using the tuple (bridge domain ID, VLAN, and destination MAC address).
4. If a match is observed for the MAC address, the result of the lookup points to the egress next-hop. The egress Packet Forwarding Engine is used to forward the packet.
5. If a miss occurs during the lookup, the flood next-hop is determined using the mesh group ID to flood the packet.

The following two significant conditions are considered in PVLAN bridging: Only a specific port to another port forwarding is permitted. A packet drop occurs on the egress interface after traversing and consuming the fabric bandwidth. To avoid traffic dropping, the decision on whether the packet needs to be dropped arrives before traversing the fabric, thereby saving the fabric bandwidth during DoS attacks. Because multiple overlapping bridge domains exist, which denotes that the same port (promiscuous or interswitch link) appears as a member in multiple bridge domains, the MAC addresses learned in one port must be visible to ports on another bridge domain. For example, a MAC address learned on a promiscuous port must be visible to both an isolated port (isolated bridge domain) and a community port (community bridge domain) on the various community bridge domains.

To resolve this problem, a shared VLAN is used for PVLAN bridging. In the shared VLAN model, all the MACs learned across all the ports are stored in the same bridge domain (primary VLAN BD) and same VLAN (primary VLAN). When the VLAN lookup is done for the packet, the PVLAN port, PVLAN bridge domain, and the PVLAN tag or ID are also used. The following processes occur with a shared VLAN methodology:

- A source MAC address (SMAC) lookup is performed to find out whether this MAC address is learned or not. If it is not a learned address, an MLP packet (route for flooding traffic to MAC learning chips) is sent to all the other Packet Forwarding Engines that are mapped with this bridging domain. In addition, an MLP packet is also sent to the host.
- A destination MAC address (DMAC) lookup using the tuple (bridge domain ID, VLAN, and destination MAC address).
- If a match is observed for the MAC address, the result of the lookup points to the egress next-hop. The egress Packet Forwarding Engine is used to forward the packet.
- If a miss occurs during the lookup, the flood next-hop is determined using the mesh group ID to flood the packet.
- If a match occurs, the group ID is derived from the VLAN lookup table and the following validation is performed to enforce primary VLAN forwarding:

```
Steps   Source  Destination      Action
Step 1     0    {*}       Permit
Step 2     {*}    0        Permit
Step 3     1    1        Drop
Step 4   X <-> Y (X > 1 and Y > 1 and X ≠ Y Drop
```

Here, {*} is a wildcard in regular expression notation referring to any value. Step 1 ensures all forwarding from promiscuous or inter switch link ports to any other port is permitted. Step 2 ensures all forwarding from any port to promiscuous or interswitch link ports is permitted. Step 3 ensures any isolated port to another isolated port is dropped. Step 4 ensures community port forwarding is permitted only within same community(X == Y) and dropped when its across community (X ff Y).

Related Documentation

-

Flow of Frames on PVLAN Ports Overview

This topic describes the manner in which traffic that enters the different PVLAN ports, such as promiscuous, isolated, and interswitch link VLANs, is processed. Sample configuration scenarios are used to describe the transmission and processing of packets.

Assume a sample deployment in which a primary VLAN named VP contains ports, p1, p2, t1, t2, i1, i2, cx1, and cx2. The port types of these configured ports are as follows:

- Promiscuous ports = p1, p2
- ISL ports = t1, t2
- Isolated ports = i1, i2
- Community VLAN = Cx
- Community ports = cx1, cx2

Bridge domains are provisioned for each of the VLANs, namely, Vp, Vi, and Vcx. Assume the bridge domains to be configured as follows:

Vp—BD_primary_Vp (ports contained are p1, t1, i1, i2, cx1, cx2)

Vi—BD_isolate_Vi (ports contained are p1, t1, *i1, *i2)

Vcx—BD_community_Vcx (ports contained are p1, t1, cx1, cx2)

The bridge domains for community, primary, and isolated VLANs are automatically created by the system internally when you configure a bridge domain with a trunk interface, access interface, or interswitch link. The bridge domains contain the same VLAN ID corresponding to the VLANs. To use bridge domains for PVLANs, you must configure the following additional attributes:

Ingress Traffic on Isolated Ports

Consider an ingress port, i1. i1 is mapped to a bridge domain named BD_isolate_Vi. BD_isolate_Vi does not have any isolated ports as an egress member. Frames can only be sent in the egress direction on p1 and t1. When a frame is sent out on p1, it is tagged with the tag of Primary VLAN Vp. A VLAN translation of Vi to Vp is performed. When a frame is propagated out of t1, it is tagged with the tag Vi.

Ingress Traffic on Community ports

Consider an ingress port as cx1. cx1 is mapped to bridge domain BD_community_Vcx. Because of the VLAN membership with the bridge domain, frames can be sent out of p1, t1, cx1, cx2. When a frame is traversed out on p1, it is tagged with tag of Primary VLAN Vp [VLAN translation]. When a frame goes out of t1, it is tagged with tag Vcx.

Ingress Traffic on Promiscuous Ports

Consider a promisicious port p1 as the ingress port. p1 is mapped to bridge domain BD_primary_Vp. Frames can go out of any member port. When a frame goes out of t1, it is tagged with tag Vp. If another promiscuous port exists, that frame is also sent out with Vp.

Ingress Traffic on Interswitch Links

With the Vlan tag Vp, assume the ingress port as t1 mapped to bridge domain BD_primary_Vp. Frames can go out of any member port. When a frame goes out of p1, it is tagged with tag Vp. With the Vlan tag Vi, t1 mapped to bridge domain BD_isolate_Vi. The frame can not egress isolated ports as they are ingress-only members of BD_isolate_Vi. When a frame goes out on p1, it is tagged with tag of Primary VLAN Vp (VLAN translation). When a frame goes out of any other trunk port, it contains the Vi tag. With the Vlan tag Vcx, t1 is mapped to BD_community_Vcx. Frames can go out of p1, t1, cx1, and cx2. When a frame goes out on p1, it is tagged with the tag of primary VLAN Vp (VLAN translation).

Packet Forwarding in PVLANs

Consider a primary VLAN with the following configuration of ports:

```
Promiscuous  P1 P2
Inter Switch Link  L1 L2
Isolated     I1 I2
Community1    C11 C12
Community2    C21 C22
```

Internally, one global BD called the primary vlan BD is created that consists of all the ports. One isolation bridge domain consisting of all isolation ports in addition the promiscuous and ISL ports and one bridge domain per community is defined consisting of community ports in addition to the promiscuous and ISL ports internally configured in the system. The bridge domains with the PVLAN ports are as follows:

```
Primary Vlan BD P1 P2 L1 L2 I1 I2 C11 C12 C21 C22
Isolated BD           I1 I2 P1 P2 L1 L2
Community1 BD         C11 C12 P1 P2 L1 L2
Community 2 BD          C21 C22 P1 P2 L1 L2
```

The following PVLAN forwarding events take place among these ports with the appropriate VLAN translation as described in the following table:

Port Type To: → From:↓	Isolated	Community	Promiscuous	Inter-switch Link
Isolated	Dropped	Dropped	Primary VLAN tag to Isolation VLAN tag.	If received with the primary VLAN tag, translate to the isolation VLAN Tag; else dropped
Promiscuous	Dropped	No translation if it is the same community; else dropped.	Primary VLAN tag to Community VLAN tag.	If received with primary VLAN tag, translate to community VLAN tag; else no translation if received with same community vlan else dropped.
Community	Isolated VLAN tag to Primary VLAN tag	Community VLAN tag to Primary VLAN tag	No translation	If received with isolation or community VLAN tag, translate to Primary VLAN tag; else no translation
Interswitch Link	No translation	No translation	No translation	No translation

Guidelines for Configuring PVLANs on MX Series Routers

Consider the following guidelines while you configure PVLANs on MX Series routers that function in enhanced LAN mode:

- PVLANs are supported on MX80 routers, on MX240, MX480, and MX960 routers with DPCs in LAN mode, on MX Series routers with MPCs.
- Isolated ports, promiscuous ports, community ports, and interswitch links (ISL) adhere to the following rules of tagging and forwarding:
  - The frames received on the primary VLAN on promiscuous ports can go to any port.
  - The frames received on isolated ports can only go to promiscuous ports and ISL ports.
  - The frames received on community ports can only go to ports of the same community, promiscuous ports, and ISL ports.
  - The frames received on ISL ports with an isolation VLAN tag or ID can only go to promiscuous ports or ISL ports.
  - The frames received on ISL ports with a community VLAN tag can only go to promiscuous ports, ISL ports, or ports belonging to a corresponding community port.
  - The frames being sent out of promiscuous ports should have a primary VLAN tag or should be untagged. It is considered untagged if the port is configured as an untagged member of the primary VLAN. The frames going out of isolated or community ports are generally untagged. However, they can also be tagged depending on the port configuration. In any case, the configured VLAN tag must be the same as the related isolated VLAN tag or community VLAN tag.
  - The frames going out of ISL ports are tagged with the primary VLAN if they are received on a promiscuous port. An untagged frame cannot exit out of an ISL port in the context of a primary VLAN, isolated VLAN, or community VLAN, but for any other VLAN, it can be untagged depending on the configuration.
  - The frames going out of ISL ports are tagged with an isolated VLAN (isolation ID) if received on the isolated port.
  - The frames going out of ISL ports are tagged with the community VLAN tag, if it is received on the corresponding community port.
- Graceful Routing Engine switchover (GRES) is supported for PVLANs.
- A virtual switch instance that contains a bridge domain associated with logical interfaces is supported.
- Aggregated Ethernet (ae) interfaces for all types of ports are supported.
- Virtual private LAN service (VPLS) instances is not supported. Integrated routing and bridging (IRB) interfaces in PVLANs are supported.
- MX Series Virtual Chassis configuration is not supported.

- MC-LAG interfaces are not supported. All ports that are associated with PVLAN bridge domains cannot be mc-ae interfaces.
- IGMP snooping is not supported. Q-in-Q tunneling is not supported.

Related Documentation

-

Configuring PVLANs on MX Series Routers in Enhanced LAN Mode

You can configure a private VLAN (PVLAN) on a single MX Series router to span multiple MX Series routers. VLANs limit broadcasts to specified users. You need to specify the interswitch link (ISL) for a PVLAN, the PVLAN port types, and secondary VLANs for the PVLAN. You must create a virtual switch routing instance with a bridge domain, and associate the interfaces with the bridge domain. You can specify the secondary VLANs as isolated or community VLANs in the bridge domain.

Before you begin configuring a PVLAN, make sure you have:

- Created and configured the necessary VLANs. See "Configuring VLAN and Extended VLAN Encapsulation" on page 256 and "Enabling VLAN Tagging" on page 247.
- Configured MX240, MX480, and MX960 routers to function in enhanced LAN mode by entering the **network-services lan** statement at the **[edit chassis]** hierarchy level.

  You must reboot the router when you configure or delete the enhanced LAN mode on the router. Configuring the **network-services lan** option implies that the system is running in the enhanced IP mode. When you configure a device to function in MX-LAN mode, only the supported configuration statements and operational show commands that are available for enabling or viewing in this mode are displayed in the CLI interface.

  If your system contains parameters that are not supported in MX-LAN mode in a configuration file, you cannot commit those unsupported attributes. You must remove the settings that are not supported and then commit the configuration. After the successful CLI commit, a system reboot is required for the attributes to become effective. Similarly, if you remove the **network-services lan** statement, the system does not run in MX-LAN mode. Therefore, all of the settings that are supported outside of the MX-LAN mode are displayed and are available for definition in the CLI interface. If your configuration file contains settings that are supported only in MX-LAN mode, you must remove those attributes before you commit the configuration. After the successful CLI commit, a system reboot is required for the CLI parameters to take effect. The Layer 2 Next-Generation CLI configuration settings are supported in MX-LAN mode. As a result, the typical format of CLI configurations might differ in MX-LAN mode.

To configure a PVLAN:

1. Create a promiscuous port for the PVLAN.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode trunk**

user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

2. Create the interswitch link (ISL) trunk port for the PVLAN.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode trunk inter-switch-link**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

3. Create the isolated port for the PVLAN. The port is identified as an isolated port or a community port, based on the VLAN ID or the list of VLAN IDs to which the interface corresponds. For example, if you configure a port with a VLAN ID of 50, and if you specify a VLAN ID of 50 as the isolated VLAN or tag in the bridge domain, the port is considered as an isolation port.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode access**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

4. Create the community port for the PVLAN. The port is identified as an isolated port or a community port, based on the VLAN ID or the list of VLAN IDs to which the interface corresponds. For example, if you configure a port with a VLAN ID of 50, and if you specify a VLAN ID of 50 as the community VLAN or tag in the bridge domain, the port is considered as a community port.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode access**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

5. Create a virtual switch instance with a bridge domain and associate the logical interfaces.

   [edit routing-instances]
   user@host# **set** *routing-instance-name* **instance-type virtual-switch**
   user@host# **set** *routing-instance-name* **interface** *interface-name* **unit** *logical-unit-number*
   user@host# **set** *routing-instance-name* **bridge-domains** *bridge-domain-name*

6. Specify the primary, isolated, and community VLAN IDs, and associate the VLANs with the bridge domain.

   [edit routing-instances *instance-name* bridge-domains *bridge-domain-name*]
   user@host# **set vlan-id** *vlan-id*
   user@host# **set isolated-vlan** *vlan-id*
   user@host# **set community-vlans** *[number number-number]*

Related Documentation

•

Example: Configuring PVLANs with Secondary VLAN Trunk Ports and Promiscuous Access Ports on a QFX Series Switch

This example shows how to configure secondary VLAN trunk ports and promiscuous access ports as part of a private VLAN configuration. Secondary VLAN trunk ports carry secondary VLAN traffic.

For a given private VLAN, a secondary VLAN trunk port can carry traffic for only one secondary VLAN. However, a secondary VLAN trunk port can carry traffic for multiple secondary VLANs as long as each secondary VLAN is a member of a different private (primary) VLAN. For example, a secondary VLAN trunk port can carry traffic for a community VLAN that is part of primary VLAN pvlan100 and also carry traffic for an isolated VLAN that is part of primary VLAN pvlan400.

To configure a trunk port to carry secondary VLAN traffic, use the *isolated* and **interface** statements, as shown in steps 12 and 13 of the example configuration for Switch 1.

NOTE: **When traffic egresses from a secondary VLAN trunk port, it normally carries the tag of the primary VLAN that the secondary port is a member of. If you want traffic that egresses from a secondary VLAN trunk port to retain its secondary VLAN tag, use the *extend-secondary-vlan-id* statement.**

A promiscuous access port carries untagged traffic and can be a member of only one primary VLAN. Traffic that ingresses on a promiscuous access port is forwarded to the ports of the secondary VLANs that are members of the primary VLAN that the promiscuous access port is a member of. This traffic carries the appropriate secondary VLAN tags when it egresses from the secondary VLAN ports if the secondary VLAN port is a trunk port.

To configure an access port to be promiscuous, use the *promiscuous* statement, as shown in step 12 of the example configuration for Switch 2.

If traffic ingresses on a secondary VLAN port and egresses on a promiscuous access port, the traffic is untagged on egress. If tagged traffic ingresses on a promiscuous access port, the traffic is discarded.

Requirements

This example uses the following hardware and software components:

- Two QFX devices
- Junos OS Release 12.2 or later for the QFX Series

Overview and Topology

Figure 24 on page 300 shows the topology used in this example. Switch 1 includes several primary and secondary private VLANs and also includes two secondary VLAN trunk ports configured to carry secondary VLANs that are members of primary VLANs pvlan100 and pvlan400.

Switch 2 includes the same private VLANs. The figure shows xe-0/0/0 on Switch 2 as configured with promiscuous access ports or promiscuous trunk ports. The example configuration included here configures this port as a promiscuous access port.

The figure also shows how traffic would flow after ingressing on the secondary VLAN trunk ports on Switch 1.

Figure 24: PVLAN Topology with Secondary VLAN Trunk Ports and Promiscuous Access Port

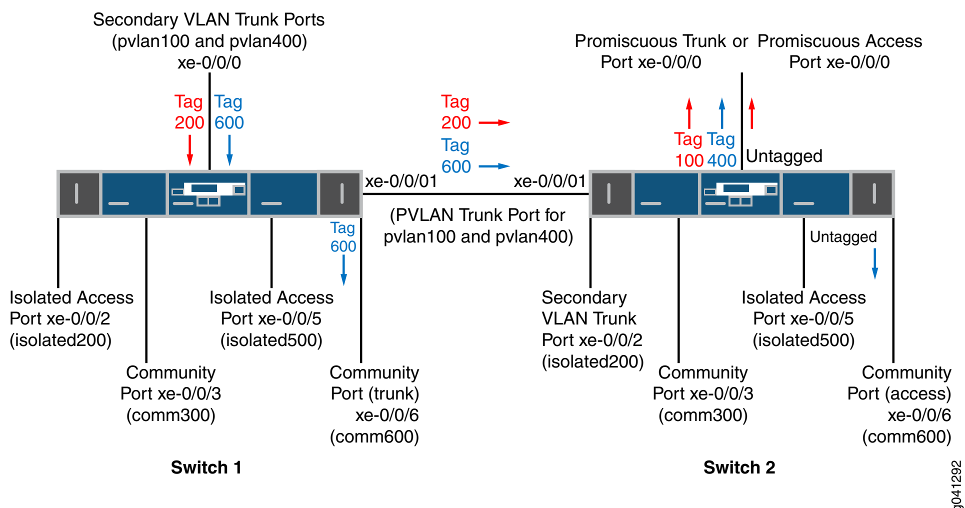

Table 22 on page 300 and Table 23 on page 301 list the settings for the example topology on both switches.

Table 22: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 1

Component	Description
pvlan100, ID 100	Primary VLAN
pvlan400, ID 400	Primary VLAN
comm300, ID 300	Community VLAN, member of pvlan100
comm600, ID 600	Community VLAN, member of pvlan400
isolation-vlan-id 200	VLAN ID for isolated VLAN, member of pvlan100

Table 22: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 1 (continued)

Component	Description
isolation–vlan-id 500	VLAN ID for isolated VLAN, member of pvlan400
xe-0/0/0.0	Secondary VLAN trunk port for primary VLANs pvlan100 and pvlan400
xe-0/0/1.0	PVLAN trunk port for primary VLANs pvlan100 and pvlan400
xe-0/0/2.0	Isolated access port for pvlan100
xe-0/0/3.0	Community access port for comm300
xe-0/0/5.0	Isolated access port for pvlan400
xe-0/0/6.0	Community trunk port for comm600

Table 23: Components of the Topology for Configuring a Secondary VLAN Trunk on Switch 2

Component	Description
pvlan100, ID 100	Primary VLAN
pvlan400, ID 400	Primary VLAN
comm300, ID 300	Community VLAN, member of pvlan100
comm600, ID 600	Community VLAN, member of pvlan400
isolation-vlan-id 200	VLAN ID for isolated VLAN, member of pvlan100
isolation–vlan-id 500	VLAN ID for isolated VLAN, member of pvlan400
xe-0/0/0.0	Promiscuous access port for primary VLANs pvlan100
xe-0/0/1.0	PVLAN trunk port for primary VLANs pvlan100 and pvlan400
xe-0/0/2.0	Secondary trunk port for isolated VLAN, member of pvlan100
xe-0/0/3.0	Community access port for comm300
xe-0/0/5.0	Isolated access port for pvlan400
xe-0/0/6.0	Community access port for comm600

Configuring the PVLANs on Switch 1

CLI Quick Configuration

To quickly create and configure the PVLANs on Switch 1, copy the following commands and paste them into a switch terminal window:

```
[edit]
```

```
set interfaces xe-0/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
set interfacesxe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400
set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/3 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/5 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode trunk
set vlans pvlan100 vlan-id 100
set vlans pvlan400 vlan-id 400
set vlans pvlan100 pvlan
set vlans pvlan400 pvlan
set vlans pvlan100 interface xe-0/0/1.0 pvlan-trunk
set vlans pvlan400 interface xe-0/0/1.0 pvlan-trunk
set vlans comm300 vlan-id 300
set vlans comm300 primary-vlan pvlan100
set vlans comm300 interface xe-0/0/3.0
set vlans comm600 vlan-id 600
set vlans comm600 primary-vlan pvlan400
set vlans comm600 interface xe-0/0/6.0
set vlans pvlan100 pvlan isolation-vlan-id 200
set vlans pvlan400 pvlan isolation-vlan-id 500
set vlans pvlan100 interface xe-0/0/0.0 isolated
set vlans pvlan400 interface xe-0/0/0.0 isolated
set vlans comm600 interface xe-0/0/0.0
set vlans pvlan100 interface xe-0/0/2.0 isolated
set vlans pvlan400 interface xe-0/0/5.0 isolated
```

Step-by-Step Procedure

To configure the private VLANs and secondary VLAN trunk ports:

1. Configure the interfaces and port modes:

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400
user@switch# set xe-0/0/2 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/3 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/5 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/6 unit 0 family ethernet-switching port-mode access
```

2. Create the primary VLANs:

```
[edit vlans]
user@switch# set pvlan100 vlan-id 100
user@switch# set pvlan400 vlan-id 400
```

NOTE: Primary VLANs must always be tagged VLANs, even if they exist on only one device.

3. Configure the primary VLANs to be private:


```
[edit vlans]
user@switch# set pvlan100 pvlan
user@switch# set pvlan400 pvlan
```

4. Configure the PVLAN trunk port to carry the private VLAN traffic between the switches:

```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/1.0 pvlan-trunk
user@switch# set pvlan400 interface xe-0/0/1.0 pvlan-trunk
```

5. Create secondary VLAN comm300 with VLAN ID 300:

```
[edit vlans]
user@switch# set comm300 vlan-id 300
```

6. Configure the primary VLAN for comm300:

```
[edit vlans]
user@switch# set comm300 primary-vlan pvlan100
```

7. Configure the interface for comm300:

```
[edit vlans]
user@switch# set comm300 interface xe-0/0/3.0
```

8. Create secondary VLAN comm600 with VLAN ID 600:

```
[edit vlans]
user@switch# set comm600 vlan-id 600
```

9. Configure the primary VLAN for comm600:

```
[edit vlans]
user@switch# set comm600 primary-vlan pvlan400
```

10. Configure the interface for comm600:

```
[edit vlans]
user@switch# set comm600 interface xe-0/0/6.0
```

11. Configure the interswitch isolated VLANs:

```
[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 200
user@switch# set pvlan400 pvlan isolation-vlan-id 500
```

NOTE: When you configure a secondary VLAN trunk port to carry an isolated VLAN, you must also configure an *isolation-vlan-id*. This is true even if the isolated VLAN exists only on one switch.

12. Enable trunk port xe-0/0/0 to carry secondary VLANs for the primary VLANs:

```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/0.0 isolated
user@switch# set pvlan400 interface xe-0/0/0.0 isolated
```

13. Configure trunk port xe-0/0/0 to carry comm600 (member of pvlan400):

```
[edit vlans]
user@switch# set comm600 interface xe-0/0/0.0
```

NOTE: You do not need to explicitly configure xe-0/0/0 to carry the isolated VLAN traffic (tags 200 and 500) because all the isolated ports in pvlan100 and pvlan400—including xe-0/0/0.0—are automatically included in the isolated VLANs created when you configured **isolation-vlan-id 200** and **isolation-vlan-id 500**.

14. Configure xe-0/0/2 and xe-0/0/6 to be isolated:

```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/2.0 isolated
user@switch# set pvlan400 interface xe-0/0/5.0 isolated
```

Results

Check the results of the configuration on Switch 1:

```
[edit]
user@switch# show
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members pvlan100;
          members pvlan400;
        }
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members pvlan100;
          members pvlan400;
        }
      }
```

```
        }
      }
      xe-0/0/2 {
        unit 0 {
          family ethernet-switching {
            port-mode access;
          }
        }
      }
      xe-0/0/3 {
        unit 0 {
          family ethernet-switching {
            port-mode access;
          }
        }
      }
      xe-0/0/5 {
        unit 0 {
          family ethernet-switching {
            port-mode access;
          }
        }
      }
      xe-0/0/6 {
        unit 0 {
          family ethernet-switching {
            port-mode trunk;
          }
        }
      }
    }
    vlans {
      comm300 {
        vlan-id 300;
          interface {
          xe-0/0/3.0;
        }
        primary-vlan pvlan100;
      }
      comm600 {
        vlan-id 600;
          interface {
          xe-0/0/6.0;
        }
        primary-vlan pvlan400;
      }
      pvlan100 {
        vlan-id 100;
        interface {
          xe-0/0/0.0;
          xe-0/0/2.0;
          xe-0/0/3.0;
          xe-0/0/1.0 {
            pvlan-trunk;
          }
        }
```

```
      no-local-switching;
      isolation-id 200;
    }
    pvlan400 {
      vlan-id 400;
      interface {
        xe-0/0/0.0;
        xe-0/0/5.0;
        xe-0/0/6.0;
        xe-0/0/1.0 {
          pvlan-trunk;
        }
      }
      no-local-switching;
      isolation-id 500;
    }
  }
```

Configuring the PVLANs on Switch 2

The configuration for Switch 2 is almost identical to the configuration for Switch 1. The most significant difference is that xe-0/0/0 on Switch 2 is configured as a promiscuous trunk port or a promiscuous access port, as Figure 24 on page 300 shows. In the following configuration, xe-0/0/0 is configured as a promiscuous access port for primary VLAN pvlan100.

If traffic ingresses on VLAN-enabled port and egresses on a promiscuous access port, the VLAN tags are dropped on egress and the traffic is untagged at that point. For example, traffic for comm600 ingresses on the secondary VLAN trunk port configured on xe-0/0/0.0 on Switch 1 and carries tag 600 as it is forwarded through the secondary VLAN. When it egresses from xe-0/0/0.0 on Switch 2, it will be untagged if you configure xe-0/0/0.0 as a promiscuous access port as shown in this example. If you instead configure xe-0/0/0.0 as a promiscuous trunk port (port-mode trunk), the traffic for comm600 carries its primary VLAN tag (400) when it egresses.

CLI Quick Configuration

To quickly create and configure the PVLANs on Switch 2, copy the following commands and paste them into a switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400
set interfaces xe-0/0/2 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/3 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/5 unit 0 family ethernet-switching port-mode access
set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode access
set vlans pvlan100 vlan-id 100
set vlans pvlan400 vlan-id 400
set vlans pvlan100 pvlan
set vlans pvlan400 pvlan
set vlans pvlan100 interface xe-0/0/1.0 pvlan-trunk
set vlans pvlan400 interface xe-0/0/1.0 pvlan-trunk
set vlans comm300 vlan-id 300
```



```
set vlans comm300 primary-vlan pvlan100
set vlans comm300 interface xe-0/0/3.0
set vlans comm600 vlan-id 600
set vlans comm600 primary-vlan pvlan400
set vlans comm600 interface xe-0/0/6.0
set vlans pvlan100 pvlan isolation-vlan-id 200
set vlans pvlan400 pvlan isolation-vlan-id 500
set vlans pvlan100 interface xe-0/0/0.0 promiscuous
set vlans pvlan100 interface xe-0/0/2.0 isolated
set vlans pvlan400 interface xe-0/0/5.0 isolated
```

Step-by-Step Procedure

To configure the private VLANs and secondary VLAN trunk ports:

1. Configure the interfaces and port modes:

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/1 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan100
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members pvlan400
user@switch# set xe-0/0/2 unit 0 family ethernet-switching port-mode trunk
user@switch# set xe-0/0/3 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/5 unit 0 family ethernet-switching port-mode access
user@switch# set xe-0/0/6 unit 0 family ethernet-switching port-mode access
```

2. Create the primary VLANs:

```
[edit vlans]
user@switch# set pvlan100 vlan-id 100
user@switch# set pvlan400 vlan-id 400
```

3. Configure the primary VLANs to be private:

```
[edit vlans]
user@switch# set pvlan100 pvlan
user@switch# set pvlan400 pvlan
```

4. Configure the PVLAN trunk port to carry the private VLAN traffic between the switches:

```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/1.0 pvlan-trunk
user@switch# set pvlan400 interface xe-0/0/1.0 pvlan-trunk
```

5. Create secondary VLAN comm300 with VLAN ID 300:

```
[edit vlans]
user@switch# set comm300 vlan-id 300
```

6. Configure the primary VLAN for comm300:

```
[edit vlans]
user@switch# set comm300 primary-vlan pvlan100
```

7. Configure the interface for comm300:

```
[edit vlans]
user@switch# set comm300 interface xe-0/0/3.0
```

8. Create secondary VLAN comm600 with VLAN ID 600:

```
[edit vlans]
user@switch# set comm600 vlan-id 600
```

9. Configure the primary VLAN for comm600:

```
[edit vlans]
user@switch# set comm600 primary-vlan pvlan400
```

10. Configure the interface for comm600:

```
[edit vlans]
user@switch# set comm600 interface xe-0/0/6.0
```

11. Configure the interswitch isolated VLANs:

```
[edit vlans]
user@switch# set pvlan100 pvlan isolation-vlan-id 200
user@switch# set pvlan400 pvlan isolation-vlan-id 500
```

12. Configure access port xe-0/0/0 to be promiscuous for pvlan100:

```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/0.0 promiscuous
```

NOTE: A promiscuous access port can be a member of only one primary VLAN.

13. Configure xe-0/0/2 and xe-0/0/6 to be isolated:

```
[edit vlans]
user@switch# set pvlan100 interface xe-0/0/2.0 isolated
user@switch# set pvlan400 interface xe-0/0/5.0 isolated
```

Results

Check the results of the configuration on Switch 2:

```
[edit]
user@switch# show
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode access;
        vlan {
```



```
                members pvlan100;
              }
            }
          }
        }
          xe-0/0/1 {
          unit 0 {
            family ethernet-switching {
              port-mode trunk;
              vlan {
                members pvlan100;
                members pvlan400;
              }
            }
          }
        }
        xe-0/0/2 {
          unit 0 {
            family ethernet-switching {
              port-mode trunk;
            }
          }
        }
        xe-0/0/3 {
          unit 0 {
            family ethernet-switching {
              port-mode access;
            }
          }
        }
        xe-0/0/5 {
          unit 0 {
            family ethernet-switching {
              port-mode access;
            }
          }
        }
        xe-0/0/6 {
          unit 0 {
            family ethernet-switching {
              port-mode access;
            }
          }
        }
      vlans {
        comm300 {
          vlan-id 300;
            interface {
            xe-0/0/3.0;
          }
          primary-vlan pvlan100;
        }
        comm600 {
          vlan-id 600;
            interface {
            xe-0/0/6.0;
```

```
        }
        primary-vlan pvlan400;
      }
      pvlan100 {
        vlan-id 100;
        interface {
          xe-0/0/0.0;
          xe-0/0/2.0;
          xe-0/0/3.0;
          xe-0/0/1.0 {
            pvlan-trunk;
          }
        }
        no-local-switching;
        isolation-id 200;
      }
      pvlan400 {
        vlan-id 400;
        interface {
          xe-0/0/5.0;
          xe-0/0/6.0;
          xe-0/0/1.0 {
            pvlan-trunk;
          }
        }
        no-local-switching;
        isolation-id 500;
      }
    }
```

Verification

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

- Verifying That the Private VLAN and Secondary VLANs Were Created on page 310
- Verifying The Ethernet Switching Table Entries on page 311

Verifying That the Private VLAN and Secondary VLANs Were Created

Purpose Verify that the primary VLAN and secondary VLANs were properly created on Switch 1.

Action Use the **show vlans** command:

```
user@switch> show vlans private-vlan

Name              Role        Tag     Interfaces
pvlan100          Primary     100     xe-0/0/0.0, xe-0/0/1.0, xe-0/0/2.0,
xe-0/0/3.0
__iso_pvlan100__  Isolated    200     xe-0/0/2.0
comm300           Community   300     xe-0/0/3.0

pvlan400          Primary     400     xe-0/0/0.0, xe-0/0/1.0, xe-0/0/5.0,
xe-0/0/6.0
__iso_pvlan400__  Isolated    500     xe-0/0/5.0
comm600           Community   600    xe-0/0/6.0
```


Meaning The output shows that the private VLANs were created and identifies the interfaces and secondary VLANs associated with them.

Verifying The Ethernet Switching Table Entries

Purpose Verify that the Ethernet switching table entries were created for primary VLAN pvlan100.

Action Show the Ethernet switching table entries for pvlan100.

```
user@switch> show ethernet-switching table vlan pvlan100 private-vlan
Ethernet-switching table: 0 unicast entries
  pvlan100              *                 Flood          - All-members
  pvlan100              00:10:94:00:00:02 Learn            xe-0/0/2.0
  __iso_pvlan100__      *                 Flood          - All-members
  __iso_pvlan100__      00:10:94:00:00:02 Replicated      - xe-0/0/2.0
```

Related Documentation

- *Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANs*
- Understanding Private VLANs on page 275
- *Understanding PVLAN Traffic Flows Across Multiple Switches*
- *Understanding Egress Firewall Filters with PVLANs*

IRB Interfaces in Private VLANs on MX Series Routers

You can configure integrated routing and bridging (IRB) interfaces in a private VLAN (PVLAN) on a single MX router to span multiple MX routers. PVLANs limit the communication within a VLAN by restricting traffic flows through their member switch ports (which are called “private ports”) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. IRB provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another bridge domain that has an IRB interface configured. You configure a logical routing interface by including the **irb** statement at the **[edit interfaces]** hierarchy level and include that interface in the bridge domain.

PVLANs are supported on MX80 routers, on MX240, MX480, and MX960 routers with DPCs in LAN mode, and on MX Series routers with MPC1, MPC2, and Adaptive Services PICs. This functionality is supported only on MX240, MX480, and MX960 routers that function in enhanced LAN mode (by entering the **network-services lan** statement at the **[edit chassis]** hierarchy level).

IRB in PVLANs replaces the external router used for routing across VLANs. The routing operations in the absence of IRB occur through external router connected to promiscuous port. This behavior takes care of all the routed frames for all the ports defined under the PVLAN domain. In this case, no layer 3 exchange occurs on MX Series routers in enhanced

LAN mode for this PVLAN bridge domain. In the case of IRB, the Layer 3 interface is associated with the primary VLAN that is configured and is considered to be a single Layer 3 interface for the entire PVLAN domain. The ingress routed traffic from all ports in the PVLAN domain needs to be mapped to this IRB interface. The egress of the IRB interface take places under the PVLAN. For a PVLAN domain spanning multiple switches, only one IRB interface can be configured in one switch. This IRB interface represent whole PVLAN domain to interact with the Layer 3 domains. An IRB interface only associates with the primary bridge domain and all Layer 3 forwarding occurs only in the primary bridge domain. When a Layer 3 packet is received in an isolated port or a promiscuous port, the device first locates the secondary bridge domain, based on secondary bridge domain to find primary bridge domain identifier. If the destination MAC address is the local IRB MAC address, the microcode transmits the packet to IRB interface associated with primary bridge domain for further processing. The same procedure occurs for receiver Layer 3 packets in an interswitch link (ISL) port with the isolated or community VLAN tag.

For the ingress Layer 3 packet with Layer 3 forwarding logic sent to IRB interfaces associated with a PVLAN bridge domain, the device processes and determines the ARP entry to send packet to the related interface that might be an isolated port or a community port. The microcode appends or translates the packet VLAN ID to the isolation or community vlan ID based on the port type. The VLAN ID is removed if the related port is untagged. A special operational case exists for Layer 3 packets that are forwarded to remote isolated or community port through the ISL link. The Layer 3 packet might contain the primary bridge domain VLAN ID and the remote node performs the translation or pop operation when it sends the packet out on the related port. This method of processing is different from Layer 2 domains. Because all forwarding base on ARP must be unicast traffic and in the remote node, the port that must be used to forward is known and the transmission of PVLAN ID occurs properly.

An ARP entry carries only the primary bridge domain information. When an ARP response is received from an isolated port or a promiscuous port, the system identifies the secondary bridge domain, and based on the secondary bridge domain, it attempts to retrieve the primary bridge domain identifier. ARP packets eventually reach the IRB interface associated with the primary bridge domain. The kernel considers this ARP packet as a normal bridge domain and creates and maintains the ARP entry only for the primary bridge domain. The same procedure is adopted for ARP request packets that are destined for the local IRB MAC address. The response is transmitted through the IRB interface and appropriate VLAN translation or a pop operation is performed, depending on the received interface.

Guidelines for Configuring IRB Interfaces in PVLANs on MX Series Routers

Keep the following points in mind when you configure IRB interfaces for PVLANs:

- All of the IP applications such as IP multicast, IPv4, IPv6, and VRRP that are compatible with IRB in normal bridge domains function properly when IRB for PVLAN bridge domains is configured.
- MC-LAG interfaces are not supported. All ports that are associated with PVLAN bridge domains cannot be mc-ae interfaces.

- IGMP snooping is not supported.
- A virtual switch instance that contains a bridge domain associated with logical interfaces is supported.
- Q-in-Q tunneling is not supported.
- Logical systems are not supported.
- Virtual private LAN service (VPLS) and Etherent VPN (EVPN) in virtual switch routing instances are not supported. A validation is performed if you attempt to configure Layer 3 interfaces in a secondary VLAN.
- MX Series Virtual Chassis configuration is not supported.

Forwarding of Packets Using IRB Interfaces in PVLANs

This topic describes how PVLAN packet forwarding operates with IRB interfaces on MX Series routers in enhanced LAN mode. The IRB interface operates as a Layer 3 gateway for all members of a bridging domain. All the members of bridging domain are assumed to be in the same subnet as the subnet of the IRB interface, which works as a gateway.

Consider a sample deployment scenario in which two routers, Router1 and Router2, are configured with a PVLAN. On Router1, the promiscuous port is P1, interswitch link is L1, isolated port is I1, and two community ports are C11 and C21. Similarly, on Router2, the promiscuous port is P2, interswitch link is L2, isolated port is I2, and two community ports are C12 and C22. In the example configuration, the two routers are interconnected through an ISL link, L1 with L2. A PVLAN domain is defined across these two routers encompassing a subdomain of isolated ports (I1, I2), and Community1 ports (C11, C12), and Community2 ports (C21, C22). Because all the ports are in the same subnet, without IRB, switching capability works across ports, across routers following the PVLAN rules. When the end-host needs to reach out cross the subnet, you must configure IRB on the bridging domain. From an end-host perspective, to reach out across the bridging domain, it needs to be configured with the IRB IP address as the default gateway address. All Layer 3 connectivity is established by processing ARP request and ARP responses. The following sections describe the different scenarios encountered for Layer 3 traffic support in PVLANs.

Incoming ARP Requests on PVLAN Ports

ARP requests enter a PVLAN port as broadcast packets. All packets that enter in the ingress direction of a PVLAN domain contain their bridge domain ID translated into the primary VLAN bridge domain ID. In this case, the bridge domain ID contained in the ARP packet is also translated be to the bridge domain ID of the primary VLAN. When IRB is configured in a bridging domain, the IRB MAC address is added to the MAC table as an eligible destination MAC address on the primary VLAN bridge domain ID. The ARP request is flooded to all ports of the secondary bridging domain in which it was received and, in addition, a copy is sent to the IRB logical interface.

When an IRB logical interface receives this packet, it sends the packet to the host as an ARP packet with the primary BD and the Layer 2 logical interface on which it is received. The PVLAN domain learns the source MAC address of the ARP packet and the kernel learns the sender IP of the ARP packet, and triggers a next-hop installation. If the ARP request is destined for IRB IP address, then an ARP response is sent. If proxy ARP is enabled on IRB, IRB responds with an ARP reply if the destination IP address is known.

The preceding configuration case describes a scenario the ARP request came on Local PVLAN port. If the ARP request is received on a remote PVLAN port, then it is flooded on all the ports of the remote PVLAN domain. Because IRB is configured only on one router of the PVLAN domain, on the remote PVLAN, the flooding is on all the ports. As part of the flooding in the remote PVLAN domain, a copy of the packet is sent to the ISL port. The ISL port processes this packet as though it was received on the local isolated port or community port and the aforementioned method of processing takes place

Outgoing ARP Responses on PVLAN Ports

When a ARP request is received in the kernel, both the bridge domain ID and the receiving Layer 2 logical interface are transmitted. A next-hop installation is triggered to create a next-hop to the Layer 2 logical interface for the sender IP address with the IRB MAC Address as the destination MAC address and the sender MAC address as the source MAC address, with both these addresses appearing as Layer 2 rewrite during the next-hop. If the ARP request queries for the IRB IP address, then an ARP response is sent to the receiving Layer 2 logical interface. If the ARP request queries for an IP address other than the IRB IP address, it is processed as though proxy ARP is enabled on IRB or it is discarded. Because all ARP requests are processed as being received on the primary VLAN, the response is also sent with the primary VLAN. However, when it reaches the receiving Layer 2 logical interface, the appropriate VLAN translation takes place.

The preceding scenario describes an ARP response being sent on a local PVLAN port. If the ARP request is received from a remote PVLAN domain, the receiving Layer 2 logical interface is the ISL port. In this case, the ARP response is sent to the ISL port, on the remote PVLAN domain, the ARP response received on the ISL port is forwarded to the same port where the ARP request is received. This behavior is possible because the source MAC address of the ARP request is learned on the shared VLAN.

Outgoing ARP Requests on PVLAN Ports

When IRB has to advertise a ARP request, it uses the kernel flood next-hop for the primary VLAN and floods to all the ports in the local PVLAN domain. The receiving ISL port also floods the packet to the remote PVLAN domain. Although the ARP request is constructed with the primary VLAN, in the egress direction, appropriate VLAN translation or VLAN pop is performed using the specific port.

Incoming ARP Responses on PVLAN Ports

ARP responses are unicast packets with the destination MAC address as the IRB MAC Address. When such a packet is received on the local PVLAN domain where IRB is enabled, it is forwarded to the IRB logical interface. When the packet arrives at the IRB logical interface, it is propagated to the host. The kernel triggers a next-hop installation with the appropriate Layer 2 rewrite. This operation works properly for ARP responses received on the local PVLAN port. If the ARP response is received on a remote PVLAN port, it is forwarded similar to a normal Layer 2 packet because IRB is not enabled in such a scenario. When the ARP request is sent out from the local PVLAN domain, the receiving ISL port in the remote PVLAN domain might have learned the IRB MAC address on that port, and this address is used to forward the packet to the IRB logical interface.

Receipt of Layer 3 Packets on PVLAN Ports

The packet is received with the IRB MAC address as the destination MAC address and it is processed through the IRB logical interface. The packet is forwarded in the same manner as a regular IP packet.

Configuring IRB Interfaces in PVLAN Bridge Domains on MX Series Routers in Enhanced LAN Mode

You can configure integrated routing and bridging (IRB) interfaces in a private VLAN (PVLAN) on a single MX router to span multiple MX routers. PVLANs limit the communication within a VLAN by restricting traffic flows through their member switch ports (which are called “private ports”) so that these ports communicate only with a specified uplink trunk port or with specified ports within the same VLAN. IRB provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another bridge domain that has an IRB interface configured. You configure a logical routing interface and include that interface in the virtual switch instance that contains the bridge domain. You can specify the secondary VLANs as isolated or community VLANs in the bridge domain.

Before you begin configuring a PVLAN, make sure you have:

- Created and configured the necessary VLANs. See “Configuring VLAN and Extended VLAN Encapsulation” on page 256 and “Enabling VLAN Tagging” on page 247.
- Configured MX240, MX480, and MX960 routers to function in enhanced LAN mode by entering the **network-services lan** statement at the **[edit chassis]** hierarchy level.

  You must reboot the router when you configure or delete the enhanced LAN mode on the router. Configuring the **network-services lan** option implies that the system is running in the enhanced IP mode. When you configure a device to function in MX-LAN mode, only the supported configuration statements and operational show commands that are available for enabling or viewing in this mode are displayed in the CLI interface.

  If your system contains parameters that are not supported in MX-LAN mode in a configuration file, you cannot commit those unsupported attributes. You must remove the settings that are not supported and then commit the configuration. After the successful CLI commit, a system reboot is required for the attributes to become effective. Similarly, if you remove the **network-services lan** statement, the system does not run in MX-LAN mode. Therefore, all of the settings that are supported outside of the MX-LAN mode are displayed and are available for definition in the CLI interface. If your configuration file contains settings that are supported only in MX-LAN mode, you must remove those attributes before you commit the configuration. After the successful CLI commit, a system reboot is required for the CLI parameters to take effect. The Layer 2 Next-Generation CLI configuration settings are supported in MX-LAN mode. As a result, the typical format of CLI configurations might differ in MX-LAN mode.

To configure an IRB interface in a PVLAN bridge domain associated with a virtual switch instance:

1. Create a promiscuous port for the PVLAN.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode trunk**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

2. Create the interswitch link (ISL) trunk port for the PVLAN.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode trunk inter-switch-link**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

3. Create the isolated port for the PVLAN. The port is identified as an isolated port or a community port, based on the VLAN ID or the list of VLAN IDs to which the interface corresponds. For example, if you configure a port with a VLAN ID of 50, and if you specify a VLAN ID of 50 as the isolated VLAN or tag in the bridge domain, the port is considered as an isolation port.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode access**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

4. Create the community port for the PVLAN. The port is identified as an isolated port or a community port, based on the VLAN ID or the list of VLAN IDs to which the interface corresponds. For example, if you configure a port with a VLAN ID of 50, and if you specify a VLAN ID of 50 as the community VLAN or tag in the bridge domain, the port is considered as a community port.

   [edit interfaces]
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge interface-mode access**
   user@host# **set interface** *interface-name* **unit** *logical-unit-number* **family bridge vlan-id** *vlan-id*

5. Create a virtual switch instance with a bridge domain and associate the logical interfaces.

   [edit routing-instances]
   user@host# **set** *routing-instance-name* **instance-type virtual-switch**
   user@host# **set** *routing-instance-name* **interface** *interface-name* **unit** *logical-unit-number*
   user@host# **set** *routing-instance-name* **bridge-domains** *bridge-domain-name*

6. Create an IRB interface and specify the IRB interface in the bridge domain associated with the virtual switch instance. IRB provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another bridge domain that has a Layer 3 protocol configured.

   [edit]
   user@host# **set interfaces irb unit *logical-unit-number* family *family-name* address *ip-address***
   [edit routing-instances *instance-name* bridge-domains *bridge-domain-name*]
   user@host# **set routing-interface irb unit *logical-unit-number***

7. Specify the primary, isolated, and community VLAN IDs, and associate the VLANs with the bridge domain.

   [edit routing-instances *instance-name* bridge-domains *bridge-domain-name*]
   user@host# **set vlan-id *vlan-id***
   user@host# **set isolated-vlan *vlan-id***
   user@host# **set community-vlans *[number number-number]***

Related Documentation
•

Example: Configuring an IRB Interface in a Private VLAN on a Single MX Series Router

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature on MX Series routers allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN.

This example describes how to create an integrated routing and bridging (IRB) interface in a PVLAN bridge domain associated with a virtual switch instance on a single MX Series router:

NOTE: **Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.**

Requirements

This example uses the following hardware and software components:

- One MX Series router in enhanced LAN mode.

- Junos OS Release 15.1 or later for MX Series routers

Before you begin configuring a PVLAN, make sure you have:

- Created and configured the necessary VLANs. See "Configuring VLAN and Extended VLAN Encapsulation" on page 256 and "Enabling VLAN Tagging" on page 247.
- Configured MX240, MX480, and MX960 routers to function in enhanced LAN mode by entering the **network-services lan** statement at the **[edit chassis]** hierarchy level.

Overview and Topology

In a large office with multiple buildings and VLANs, you might need to isolate some workgroups or other endpoints for security reasons or to partition the broadcast domain. This configuration example shows a simple topology to illustrate how to create a PVLAN with one primary VLAN and four community VLANs, as well as two isolated ports.

Assume a sample deployment in which a primary VLAN named VP contains ports, p1, p2, t1, t2, i1, i2, cx1, and cx2. The port types of these configured ports are as follows:

- Promiscuous ports = p1, p2
- ISL ports = t1, t2
- Isolated ports = i1, i2
- Community VLAN = Cx
- Community ports = cx1, cx2

An IRB interface, irb.0, is configured and mapped to the bridge domain in the virtual switch instance.

Bridge domains are provisioned for each of the VLANs, namely, Vp, Vi, and Vcx. Assume the bridge domains to be configured as follows:

Vp—BD_primary_Vp (ports contained are p1, t1, i1, i2, cx1, cx2)

Vi—BD_isolate_Vi (ports contained are p1, t1, *i1, *i2)

Vcx—BD_community_Vcx (ports contained are p1, t1, cx1, cx2)

The bridge domains for community, primary, and isolated VLANs are automatically created by the system internally when you configure a bridge domain with a trunk interface, access interface, or interswitch link. The bridge domains contain the same VLAN ID corresponding to the VLANs. To use bridge domains for PVLANs, you must configure the following additional attributes:

Configuration

To configure an IRB interface in a PVLAN, perform these tasks:

CLI Quick Configuration

To quickly create and configure a PVLAN and include an IRB interface in a PVLAN bridge domain associated with a virtual switch instance, copy the following commands and paste them into the router terminal window:

Configuring an IRB Interface

```
set interfaces irb unit 0 family inet address 22.22.22.1/24
```

Configuring Promiscuous, ISL, Isolated, and Community Ports

```
set interfaces ge-0/0/9 unit 0 family bridge interface-mode trunk
set interfaces ge-0/0/9 unit 0 family bridge vlan-id 100
set interfaces ge-0/0/13 unit 0 family bridge interface-mode trunk
set interfaces ge-0/0/13 unit 0 family bridge vlan-id 100
set interfaces ge-0/0/10 unit 0 family bridge interface-mode access
set interfaces ge-0/0/10 unit 0 family bridge vlan-id 10
set interfaces ge-0/0/12 unit 0 family bridge interface-mode access
set interfaces ge-0/0/12 unit 0 family bridge vlan-id 10
set interfaces ge-0/0/1 unit 0 family bridge interface-mode access
set interfaces ge-0/0/1 unit 0 family bridge vlan-id 50
set interfaces ge-0/0/2 unit 0 family bridge interface-mode access
set interfaces ge-0/0/2 unit 0 family bridge vlan-id 50
set interfaces ge-0/0/3 unit 0 family bridge interface-mode access
set interfaces ge-0/0/3 unit 0 family bridge vlan-id 60
set interfaces ge-0/0/4 unit 0 family bridge interface-mode access
set interfaces ge-0/0/4 unit 0 family bridge vlan-id 60
```

Configuring a Virtual Switch Instance With Bridge Domain Interfaces

```
set routing-instances vs-1 instance-type virtual-switch
set routing-instances vs-1 interface ge-0/0/1.0
set routing-instances vs-1 interface ge-0/0/2.0
set routing-instances vs-1 interface ge-0/0/3.0
set routing-instances vs-1 interface ge-0/0/4.0
set routing-instances vs-1 interface ge-0/0/9.0
set routing-instances vs-1 interface ge-0/0/10.0
set routing-instances vs-1 interface ge-0/0/12.0
set routing-instances vs-1 interface ge-0/0/13.0
set routing-instances vs-1 bridge-domains bd1
```

Specify the IRB Interface and Primary, Isolated, and Community VLAN IDs in the Bridge Domain

```
set routing-instances vs1 bridge-domains bd1 vlan-id 100
set routing-instances vs1 bridge-domains bd1 isolated-vlan 10
set routing-instances vs1 bridge-domains bd1 community-vlans [50 60]
set routing-instances vs1 bridge-domains bd1 routing-interface irb.0
```

Step-by-Step Procedure

To configure the interswitch link (ISL) for a PVLAN, the PVLAN port types, and secondary VLANs for the PVLAN:

1. Create an IRB interface.

```
[edit interfaces]
user@host# set interfaces irb unit 0 family inet address 22.22.22.1/24
```

2. Create a promiscuous port for the PVLAN.

```
[edit interfaces]
user@host# set ge-0/0/9 unit 0 family bridge interface-mode trunk
user@host# set ge-0/0/9 unit 0 family bridge vlan-id 100
```

3. Create the interswitch link (ISL) trunk port for the PVLAN.

```
[edit interfaces]
user@host# set ge-0/0/13 unit 0 family bridge interface-mode trunk inter-switch-link
user@host# set ge-0/0/13 unit 0 family bridge vlan-id 100
```

4. Create the isolated ports for the PVLAN.

```
[edit interfaces]
user@host# set ge-0/0/10 unit 0 family bridge interface-mode access
user@host# set ge-0/0/10 unit 0 family bridge vlan-id 10
user@host# set ge-0/0/12 unit 0 family bridge interface-mode access
user@host# set ge-0/0/12 unit 0 family bridge vlan-id 10
```

5. Create the community ports for the PVLAN.

```
[edit interfaces]
user@host# set ge-0/0/1 unit 0 family bridge interface-mode access
user@host# set ge-0/0/1 unit 0 family bridge vlan-id 50
user@host# set ge-0/0/2 unit 0 family bridge interface-mode access
user@host# set ge-0/0/2 unit 0 family bridge vlan-id 50
user@host# set ge-0/0/3 unit 0 family bridge interface-mode access
user@host# set ge-0/0/3 unit 0 family bridge vlan-id 60
user@host# set ge-0/0/4 unit 0 family bridge interface-mode access
user@host# set ge-0/0/4 unit 0 family bridge vlan-id 60
```

6. Create a virtual switch instance with a bridge domain and associate the logical interfaces.

```
[edit routing-instances]
user@host# set vs-1 instance-type virtual-switch
user@host# set vs-1 interface ge-0/0/1.0
user@host# set vs-1 interface ge-0/0/2.0
user@host# set vs-1 interface ge-0/0/3.0
user@host# set vs-1 interface ge-0/0/4.0
user@host# set vs-1 interface ge-0/0/9.0
user@host# set vs-1 interface ge-0/0/10.0
user@host# set vs-1 interface ge-0/0/12.0
user@host# set vs-1 interface ge-0/0/13.0
user@host# set vs-1 bridge-domains bd1
```

7. Specify the IRB interface, primary, isolated, and community VLAN IDs, and associate the VLANs with the bridge domain.

```
[edit routing-instances vs1 bridge-domains bd1]
user@host# set vlan-id 100
user@host# set isolated-vlan 10
user@host# set community-vlans [50 60]
user@host# set routing-interface irb.0
```

Results

Check the results of the configuration:

```
[edit]
[interfaces]
    ge-0/0/9 {
        unit 0 {
            family bridge {
  interface-mode trunk;
      vlan-id 100;            Promiscuous port by vlan id
            }
        }
    }

    ge-0/0/13 {
        unit 0 {
            family bridge {
  interface-mode trunk inter-switch-link;   ISL trunk
  vlan-id 100;
            }
        }
    }


    ge-0/0/10 {
        unit 0 {
            family bridge {
                  interface-mode access;
                  vlan-id 10;             isolated port by vlan ID
             }
        }
    }

    ge-0/0/12 {
        unit 0 {
            family bridge {
                  interface-mode access;
                  vlan-id 10;             isolated port by vlan ID
             }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family bridge {
                  interface-mode access;
                  vlan-id 50;             community port by vlan ID
             }
        }
    }

    ge-0/0/2 {
        unit 0 {
            family bridge {
                  interface-mode access;
                  vlan-id 50;             community port by vlan ID
             }

        }
    }
```

```
     ge-0/0/3 {
         unit 0 {
             family bridge {
                   interface-mode access;
                   vlan-id 60;              community port by vlan ID
              }
         }
     }

     ge-0/0/4 {
         unit 0 {
             family bridge {
                   interface-mode access;
                   vlan-id 60;              community port by vlan ID
              }
         }
     }

     irb {
        unit 0 {
           family inet {
             address 22.22.22.1/24;
           }
     }
}

[edit]
routing-instances {
    vs-1 {
        instance-type virtual-switch;
        interface ge-0/0/1.0;
        interface ge-0/0/2.0;
        interface ge-0/0/3.0;
        interface ge-0/0/4.0;
        interface ge-0/0/9.0;
        interface ge-0/0/10.0;
        interface ge-0/0/12.0;
        interface ge-0/0/13.0;

 bridge-domains {
     bd1 {
             vlan-id 100;                /* primary vlan */
                   isolated-vlan 10;
                   community-vlans [50 60]
                   routing-interface irb.0 /* IRB interface */
      }
}
}
```

Verification

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

- Verifying That the Private VLAN and Secondary VLANs Were Created on page 323

Verifying That the Private VLAN and Secondary VLANs Were Created

Purpose Verify that the primary VLAN and secondary VLANs were properly created on the switch.

Action Use the **show bridge domain** command:

```
user@host> show bridge domain
Routing instance          Bridge domain            VLAN ID     Interfaces
default-switch                   bd1-primary-100          100
                                                                     ge-0/0/9.0
                                          ge-0/0/10.0
                                                                    ge-0/0/12.0
                                                                    ge-0/0/13.0
                                                                     ge-0/0/1.0
                                                                     ge-0/0/2.0
                                                                     ge-0/0/3.0

                                                                     ge-0/0/4.0
default-switch                 bd1-isolation-10          10
                                                                       ge-0/0/9.0
                                                                        ge-0/0/10.0
                                                                      ge-0/0/12.0
                                                                      ge-0/0/13.0


default-switch                   bd1-comunity-50          50
                                                                       ge-0/0/9.0
                                                                        ge-0/0/13.0
                                                                        ge-0/0/1.0
                                                                       ge-0/0/2.0

default-switch                    bd1-comunity-60          60
                                                                       ge-0/0/9.0
                                                                        ge-0/0/13.0
                                                                     ge-0/0/3.0
                                                                      ge-0/0/4.0
```

Meaning The output shows that the primary VLAN was created and identifies the interfaces and secondary VLANs associated with it.

Related Documentation

-

CHAPTER 13

Configuring Layer 2 Bridging Interfaces

Layer 2 Bridging Interfaces Overview

Bridging operates at Layer 2 of the OSI reference model while routing operates at Layer 3. A set of logical ports configured for bridging can be said to constitute a bridging domain.

A bridging domain can be created by configuring a routing instance and specifying the instance-type as **bridge**.

Integrated routing and bridging (IRB) is the ability to:

- Route a packet if the destination MAC address is the MAC address of the router and the packet **ethertype** is IPv4, IPv6, or MPLS.
- Switch all multicast and broadcast packets within a bridging domain at layer 2.
- Route a copy of the packet if the destination MAC address is a multicast address and the **ethertype** is IPv4 or IPv6.
- Switch all other unicast packets at Layer 2.
- Handle supported Layer 2 control packets such as STP and LACP.
- Handle supported Layer 3 control packets such as OSPF and RIP.

Related Documentation

- Configuring Layer 2 Bridging Interfaces on page 326
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Layer 2 Bridging Interfaces

Integrated routing and bridging interfaces are logical Layer 3 VLAN interfaces that route traffic between bridge domains (VLANs). So, an IRB logical interface is usually associated with a bridge domain or VLAN. The IRB logical interface also functions as the gateway IP address for the other devices on the same sub-network that are associated with the same VLAN. IRB interfaces support Layer 2 bridging and Layer 3 routing on the same interface. As a result, IRB interfaces enable the router to act both as a router and as a Layer 2 switch at the same time.

NOTE: **If the status of all Layer 2 logical interfaces in the bridge domain is down, the status of the irb logical interface is also down.**

To configure an IRB logical interface:

1. In configuration mode, at the [**edit bridge-domains**] hierarchy level, configure the bridge domain by specifying the name of the bridge and the VLAN ID.

   [edit bridge-domains]
   user@host# set *bridge-domain-name* vlan-id *vlan-id*

2. Configure an interface in trunk mode and include the interface in the appropriate bridge domain using the **vlan-id-list** command at the [**edit interfaces**] hierarchy level.

   [edit interfaces]
   user@host# set *interfacetype-fpc/pic/port* vlan-tagging
   user@host# set *interfacetype-fpc/pic/port* unit *logical-unit-number* family bridge interface-mode trunk
   user@host# set *interfacetype-fpc/pic/port* unit *logical-unit-number* family bridge vlan-id-list *vlan-id*

3. Configure the IRB interface at the [**edit interfaces**] hierarchy level and specify the associated IP address.

   [edit interfaces]
   user@host# set interfaces irb unit *logical-unit-number* family inet address *address*

4. Configure the IRB interface as the routing interface for the bridge domain at the [**edit bridge-domains**] hierarchy level.

   [edit bridge-domains]
   user@host# set *bridge-domain- name* vlan-id *vlan-id* routing-interface irb.*logical-interface-number*

Related Documentation

- Layer 2 Bridging Interfaces Overview on page 325
- Example: Configuring the MAC Address of an IRB Interface on page 327

Example: Configuring the MAC Address of an IRB Interface

This example shows how to configure the media access control (MAC) address of an integrated routing and bridging (IRB) interface for devices with Modular Port Concentrator (MPC) cards . An IRB interface is a Layer 3 routing interface that is used in a bridge domain or virtual private LAN service (VPLS) routing.

Requirements

This example requires the following hardware and software components:

- MX Series routers with MPC cards.
- Junos OS Release 13.2 or later running on all devices.

Overview

Junos OS Release 13.2 and later support the assignment of MAC addresses to IRB logical interfaces. The IRB logical interfaces provide support for simultaneous Layer 2 bridging and Layer 3 routing within the same bridge domain. Packets that arrive on an interface of the bridge domain are either switched or routed, based on the destination MAC address of the packet. The packets with the router's Layer 2 virtual MAC address, which is manually configured, are switched to Layer 2 interfaces.

Configuring a MAC address of an IRB logical interface allows the use of a transparent firewall between two VLANs on the same switch. When both VLANs are on the same subnet and traffic from one VLAN needs to go through the firewall to the host on the other VLAN, then the VLAN tag is changed to communicate with the host on the other VLAN.

Before the introduction of this feature, if the MAC address of an IRB logical interface was the same for both VLANs, the firewall dropped the traffic. This new feature allows you to configure distinct MAC addresses for different VLANs, which facilitates the exchange of traffic between two VLANs on the same switch.

In case of VPLS multihoming, if there is a failover of the primary provider edge (PE) router to a secondary PE router, the MAC address of an IRB changes. The hosts connected to the customer edge (CE) router must change their Address Resolution Protocol (ARP) for IRB's IP and MAC address. This feature allows you to configure the same MAC address for IRB interfaces in both the primary and secondary PE routers and eliminates the need for changing the ARP binding of the IRB logical interface in CE routers, in case of a failover.

Figure 25 on page 328 shows the sample topology.

Figure 25: Configuring the MAC Address of an IRB Interface

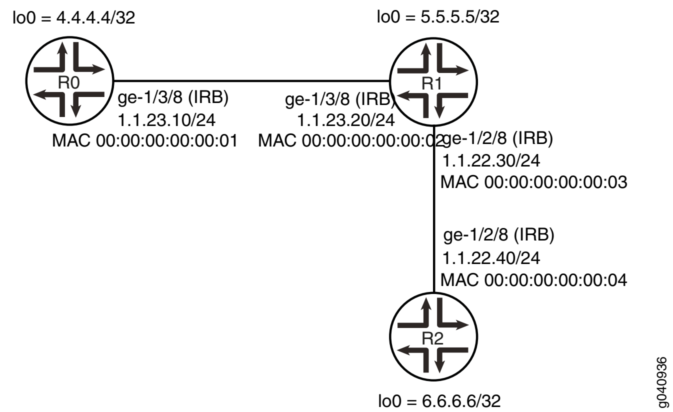

In this example you configure MAC address of IRB logical interfaces.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R0

```
set interfaces ge-1/3/8 vlan-tagging
set interfaces ge-1/3/8 encapsulation flexible-ethernet-services
set interfaces ge-1/3/8 unit 10 encapsulation vlan-bridge
set interfaces ge-1/3/8 unit 10 vlan-id 10
set interfaces irb unit 10 family inet address 1.1.23.1/24
set interfaces irb unit 10 family mpls
set interfaces irb unit 10 mac 00:00:00:00:00:01
set interfaces lo0 unit 10 family inet address 4.4.4.4/32
set protocols rsvp interface irb.10
set protocols mpls label-switched-path R0-1-R2 to 6.6.6.6
set protocols mpls label-switched-path R0-1-R2 install 6.6.6.6/32 active
set protocols mpls label-switched-path R0-1-R2 no-cspf
set protocols mpls interface irb.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 4.4.4.4
set protocols bgp group ibgp neighbor 6.6.6.6
set protocols ospf area 0.0.0.0 interface irb.10
set protocols ospf area 0.0.0.0 interface lo0.10 passive
set protocols ldp interface irb.10
set protocols ldp interface lo0.10
set routing-options autonomous-system 400
set bridge-domains lsbd1 vlan-id 10
```



```
set bridge-domains lsbd1 interface ge-1/3/8.10
set bridge-domains lsbd1 routing-interface irb.10
```

Router R1

```
set interfaces ge-1/3/8 vlan-tagging
set interfaces ge-1/3/8 encapsulation flexible-ethernet-services
set interfaces ge-1/3/8 unit 10 encapsulation vlan-bridge
set interfaces ge-1/3/8 unit 10 vlan-id 10
set interfaces ge-1/2/8 vlan-tagging
set interfaces ge-1/2/8 encapsulation flexible-ethernet-services
set interfaces ge-1/2/8 unit 40 encapsulation vlan-bridge
set interfaces ge-1/2/8 unit 40 vlan-id 40
set interfaces irb unit 20 family inet address 1.1.23.2/24
set interfaces irb unit 20 family mpls
set interfaces irb unit 20 mac 00:00:00:00:00:02
set interfaces irb unit 30 family inet address 1.1.22.2/24
set interfaces irb unit 30 family mpls
set interfaces irb unit 30 mac 00:00:00:00:00:03
set interfaces lo0 unit 20 family inet address 5.5.5.5/32
set protocols rsvp interface irb.20
set protocols rsvp interface irb.30
set protocols mpls interface irb.30
set protocols mpls interface irb.20
set protocols ospf area 0.0.0.0 interface irb.20
set protocols ospf area 0.0.0.0 interface irb.30
set protocols ospf area 0.0.0.0 interface lo0.20 passive
set protocols ldp interface irb.20
set protocols ldp interface irb.30
set protocols ldp interface lo0.20
set routing-options autonomous-system 400
set bridge-domains lsbd2 vlan-id 10
set bridge-domains lsbd2 interface ge-1/3/8.10
set bridge-domains lsbd2 routing-interface irb.20
set bridge-domains lsbd3 vlan-id 40
set bridge-domains lsbd3 interface ge-1/2/8.40
set bridge-domains lsbd3 routing-interface irb.30
```

Router R2

```
set interfaces ge-1/2/8 vlan-tagging
set interfaces ge-1/2/8 encapsulation flexible-ethernet-services
set interfaces ge-1/2/8 unit 40 encapsulation vlan-bridge
set interfaces ge-1/2/8 unit 40 vlan-id 40
set interfaces irb unit 40 family inet address 1.1.22.1/24
set interfaces irb unit 40 family mpls
set interfaces irb unit 40 mac 00:00:00:00:00:04
set interfaces lo0 unit 30 family inet address 6.6.6.6/32
set protocols rsvp interface irb.40
set protocols mpls label-switched-path R2-1-R0 to 4.4.4.4
set protocols mpls label-switched-path R2-1-R0 no-cspf
set protocols mpls interface irb.40
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 6.6.6.6
set protocols bgp group ibgp neighbor 4.4.4.4
set protocols ospf area 0.0.0.0 interface irb.40
set protocols ospf area 0.0.0.0 interface lo0.30 passive
set protocols ldp interface irb.40
```

set protocols ldp interface lo0.30
set routing-options autonomous-system 400
set bridge-domains lsbd4 vlan-id 40
set bridge-domains lsbd4 interface ge-1/2/8.40
set bridge-domains lsbd4 routing-interface irb.40

Configuring the MAC Address of an IRB Interface

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

NOTE: Repeat this procedure for Juniper Networks Routers R1 and R2, modifying the appropriate interface names, addresses, and any other parameters for each router.

To configure the MAC address of an IRB interface on Router R0:

1. Configure the physical interfaces.

```
[edit interfaces ge-1/3/8]
user@R0# set vlan-tagging
user@R0# set encapsulation flexible-ethernet-services
user@R0# set unit 10 encapsulation vlan-bridge
user@R0# set unit 10 vlan-id 10
```

2. Configure the IRB logical interface.

```
[edit interfaces irb]
user@R0# set unit 10 family inet address 1.1.23.1/24
user@R0# set unit 10 family mpls
user@R0# set unit 10 mac 00:00:00:00:00:01

[edit interfaces]
user@R0# set lo0 unit 10 family inet address 4.4.4.4/32
```

3. Configure the RSVP protocol.

```
[edit protocols rsvp]
user@R0# set interface irb.10
```

4. Configure the MPLS protocol.

```
[edit protocols mpls]
user@R0# set label-switched-path R0-1-R2 to 6.6.6.6
user@R0# set label-switched-path R0-1-R2 install 6.6.6.6/32 active
user@R0# set label-switched-path R0-1-R2 no-cspf
user@R0# set interface irb.10
user@R0# set interface irb.10
```


5. Configure the BGP protocol.

```
[edit protocols BGP]
user@R0# set group ibgp type internal
user@R0# set group ibgp local-address 4.4.4.4
user@R0# set group ibgp neighbor 6.6.6.6
```

6. Configure the OSPF protocol.

```
[edit protocols ospf]
user@R0# set area 0.0.0.0 interface irb.10
user@R0# set area 0.0.0.0 interface lo0.10 passive
```

7. Configure the LDP protocol.

```
[edit protocols ldp]
user@R0# set interface irb.10
user@R0# set interface lo0.10
```

8. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R0# set autonomous-system 400
```

9. Configure the bridge domains.

```
[edit]
user@R0# set bridge-domains lsbd1 vlan-id 10
user@R0# set bridge-domains lsbd1 interface ge-1/3/8.10
user@R0# set bridge-domains lsbd1 routing-interface irb.10
```

Results

From configuration mode, enter the **show interfaces**, **show protocols** and **show bridge-domains**, commands and confirm your configuration. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
ge-1/3/8 {
   unit 10 {
      encapsulation vlan-bridge;
      vlan-id 10;
   }
}
irb {
   unit 10 {
      family inet {
         mtu 1450;
         address 1.1.1.1/24;
         address 1.1.23.1/24;
      }
```

```
            family mpls;
            mac 00:00:00:00:00:01;
        }
    }
    lo0 {
        unit 10 {
            family inet {
                address 4.4.4.4/32;
            }
        }
    }
    user@R0# show protocols
    rsvp {
        interface irb.10;
    }
        mpls {
            label-switched-path R0-1-R2 {
                to 6.6.6.6;
                install 6.6.6.6/32 active;
                no-cspf;
            }
            interface irb.10;
        }
        bgp {
            group ibgp {
                type internal;
                local-address 4.4.4.4;
                neighbor 6.6.6.6;
            }
        }
        ospf {
            area 0.0.0.0 {
                interface irb.10;
                interface lo0.10 {
                    passive;
                }
            }
        }
        ldp {
            interface irb.10;
            interface lo0.10;
        }
    user@R0# show bridge-domains
    lsbd1 {
        vlan-id 10;
        interface ge-1/3/8.10;
        routing-interface irb.10;
    }
```

If you are done configuring the devices, commit the configuration.

```
user@host# commit
```


Verification

Confirm that the configuration is working properly.

- Verifying the MAC Address of the IRB Interface on page 333

Verifying the MAC Address of the IRB Interface

Purpose Verify that the specified MAC address is assigned to the IRB interface.

Action From operational mode, run the **show interfaces irb** command on the device.

user@host# **show interfaces irb**

```
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 505
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 80:71:1f:c2:58:f0, Hardware address: 80:71:1f:c2:58:f0
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.10 (Index 326) (SNMP ifIndex 634)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    MAC: 00:00:00:00:00:01
    Bandwidth: 1000mbps
    Routing Instance: LS1/default Bridging Domain: lsbd1+10
    Input packets : 55202
    Output packets: 69286
    Protocol inet, MTU: 1450
      Flags: Sendbcast-pkt-to-re, Is-Primary, User-MTU
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
      Addresses, Flags: Is-Preferred
        Destination: 1.1.23/24, Local: 1.1.23.1, Broadcast: 1.1.23.255
    Protocol mpls, MTU: 1500, Maximum labels: 3
      Flags: Is-Primary
    Protocol multiservice, MTU: 1500

  Logical interface irb.20 (Index 358) (SNMP ifIndex 635)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    MAC: 00:00:00:00:00:02
    Bandwidth: 1000mbps
    Routing Instance: LS2/default Bridging Domain: lsbd2+10
    Input packets : 66044
    Output packets: 68464
    Protocol inet, MTU: 1450
      Flags: Sendbcast-pkt-to-re, Is-Primary, User-MTU
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255
      Addresses, Flags: Is-Preferred
        Destination: 1.1.23/24, Local: 1.1.23.2, Broadcast: 1.1.23.255
    Protocol mpls, MTU: 1500, Maximum labels: 3
      Flags: Is-Primary
    Protocol multiservice, MTU: 1500

  Logical interface irb.30 (Index 360) (SNMP ifIndex 636)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    MAC: 00:00:00:00:00:03
    Bandwidth: 1000mbps
    Routing Instance: LS2/default Bridging Domain: lsbd3+40
    Input packets : 26948
    Output packets: 53605
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
```



```
        Destination: 1.1.22/24, Local: 1.1.22.2, Broadcast: 1.1.22.255
      Addresses, Flags: Is-Preferred
        Destination: 2.2.2/24, Local: 2.2.2.1, Broadcast: 2.2.2.255
    Protocol mpls, MTU: 1500, Maximum labels: 3
    Protocol multiservice, MTU: 1500

  Logical interface irb.40 (Index 355) (SNMP ifIndex 632)
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    MAC: 00:00:00:00:00:04
    Bandwidth: 1000mbps
    Routing Instance: LS3/default Bridging Domain: lsbd4+40
    Input packets : 40575
    Output packets: 31128
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re, Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.22/24, Local: 1.1.22.1, Broadcast: 1.1.22.255
    Protocol mpls, MTU: 1500, Maximum labels: 3
      Flags: Is-Primary
    Protocol multiservice, MTU: 1500
```

Meaning The output shows the manually configured MAC address in the MAC field.

NOTE: If you did not configure the MAC address for a logical interface, the output does not include this value. However, the device uses the MAC address of the physical interface during data transmission.

Related Documentation

- mac on page 1242
- *Active-Active Bridging and VRRP over IRB Functionality Overview*

CHAPTER 14

Configuring Link Layer Discovery Protocol

LLDP Overview

The Link Layer Discovery Protocol (LLDP) is an industry-standard, vendor-neutral method to allow networked devices to advertise capabilities, identity, and other information onto a LAN. The Layer 2 protocol, detailed in IEEE 802.1AB-2005, replaces several proprietary protocols implemented by individual vendors for their equipment.

LLDP allows network devices that operate at the lower layers of a protocol stack (such as Layer 2 bridges and switches) to learn some of the capabilities and characteristics of LAN devices available to higher layer protocols, such as IP addresses. The information gathered through LLDP operation is stored in a network device and is queried with SNMP. Topology information can also be gathered from this database.

Some of the information that can be gathered by LLDP (only minimal information is mandatory) is:

- System name and description
- Port name and description
- VLAN name and identifier
- IP network management address
- Capabilities of the device (for example, switch, router, or server)
- MAC address and physical layer information
- Power information
- Link aggregation information

NOTE: LLDP media endpoint discovery (LLDP-MED) is not supported on T Series routers.

LLDP frames are sent at fixed intervals on each port that runs LLDP. LLDP protocol data units (LLDP PDUs) are sent inside Ethernet frames and identified by their destination Media Access Control (MAC) address (**01:80:C2:00:00:0E**) and Ethertype (**0x88CC**). Mandatory information supplied by LLDP is chassis ID, port ID, and a time-to-live value for this information.

LLDP is a powerful way to allow Layer 2 devices to gather details about other network-attached devices.

Related Documentation

- Configuring LLDP on page 338
- *Obsolete: Configuring LLDP in ACX Series*
- Tracing LLDP Operations on page 344
- Example: Configuring LLDP on page 342
- LLDP Operational Mode Commands on page 343

Configuring LLDP

You configure LLDP by including the **lldp** statement and associated parameters at the **[edit protocols]** hierarchy level. The complete set of LLDP statements follows:

```
lldp {
  advertisement-interval seconds;
  disable;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  management-address ip-management-address;;
  port-description-type {
    interface-alias;
    interface-description;
  }
  port-id-subtype {
    interface-name;
    locally-assigned;
  }
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  transmit-delay seconds
}
```


The following statements have default values:

- **advertisement-interval**—The default value is 30 seconds. The allowable range is from 5 through 32768 seconds.
- **hold-multiplier**—The default values is 4. The allowable range is from 2 through 10.
- **ptopo-configuration-maximum-hold-time**—The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds.
- **transmit-delay**—The default values is 2 seconds. The allowable range is from 1 through 8192 seconds.

The following statements must be explicitly configured:

- **lldp-configuration-notification-interval**—The allowable range is from 0 through 3600 seconds. There is no default value.
- **ptopo-configuration-trap-interval**—The allowable range is from 1 through 2147483647 seconds. There is no default value.

LLDP is enabled on all interfaces by default. If it is disabled, you can enable LLDP by configuring it on all interfaces or on specific interfaces.

NOTE: **The interface-name must be the physical interface (for example, ge-1/0/0) and not a logical interface (unit).**

- To configure LLDP on all interfaces:

  ```
  [edit protocols lldp]
  user@switch# set interface all
  ```

- To configure LLDP on a specific interface:

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

To disable LLDP, include the **disable** option:

- To disable LLDP on all interfaces:

  ```
  [edit protocols lldp]
  user@switch# set interface all disable
  ```

- To disable LLDP on a specific interface:

  ```
  [edit protocols lldp]
  user@switch# set interface interface-name disable
  ```

Starting with Junos OS Release 14.2, you can configure management interfaces, such as fxp0 or me0, on MX Series routers to send LLDP frames to and receive LLDP frames from neighboring LLDP interfaces. To configure the management interfaces, nclude the **interface** ***interface-name*** statement at the **[edit protocols lldp]** and **[edit routing-instances**

routing-instance-name **protocols lldp]** hierarchy levels. By default, the functionality to send LLDP frames is enabled. You can also specify a management interface with the **show lldp neighbors interface** ***interface-name*** command to view configuration details about LLDP neighbors for the corresponding management interface.

To configure LLDP on a T Series router within a TX Matrix, you must specify the interface name in the LLDP configuration for the TX Matrix. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

Starting with Junos OS Release 14.2, LLDP is supported on extended ports in the Junos Fusion technology.For information about interface names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

The advertisement interval determines the frequency that an LLDP interface sends LLDP advertisement frames. The default value is 30 seconds. The allowable range is from 5 through 32768 seconds. You adjust this parameter by including the **advertisement-interval** statement at the **[edit protocols lldp]** hierarchy level.

The hold multiplier determines the multiplier to apply to the advertisement interval. The resulting value in seconds is used to cache learned LLDP information before discard. The default value is 4. When used with the default advertisement interval value of 30 seconds, this makes the default cache lifetime 120 seconds. The allowable range of the hold multiplier is from 2 through 10. You adjust this parameter by including the **hold-multiplier** statement at the **[edit protocols lldp]** hierarchy level.

The transmit delay determines the delay between any two consecutive LLDP advertisement frames. The default value is 2 seconds. The allowable range is from 1 through 8192 seconds. You adjust this parameter by including the **transmit-delay** statement at the **[edit protocols lldp]** hierarchy level.

The physical topology configuration maximum hold time determines the time interval for which an agent device maintains physical topology database entries. The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds. You adjust this parameter by including the **ptopo-configuration-maximum-hold-time** statement at the **[edit protocols lldp]** hierarchy level.

The LLDP configuration notification interval determines the period for which trap notifications are sent to the SNMP Master Agent when changes occur in the database of LLDP information. This capability is disabled by default. The allowable range is from 0 (disabled) through 3600 seconds. You adjust this parameter by including the **lldp-configuration-notification-interval** statement at the **[edit protocols lldp]** hierarchy level.

The physical topology configuration trap interval determines the period for which trap notifications are sent to the SNMP Master Agent when changes occur in the global physical topology statistics. This capability is disabled by default. The allowable range is from 0 (disabled) through 3600 seconds. The LLDP agent sends traps to the SNMP Master Agent if this interval has a value greater than 0 and there is any change during the **lldp-configuration-notification-interval** trap interval. You adjust this parameter by including the **ptopo-configuration-trap-interval** statement at the **[edit protocols lldp]** hierarchy level.

By default, LLDP generates the SNMP index of the interface for the port ID Type, Length, and Value (TLV). Starting with Junos OS Release 12.3R1, you can generate the interface name as the port ID TLV. To do so, include the **interface-name** statement at the **[edit protocols lldp** port-id-subtype] hierarchy level. When the **interface-name** statement is configured on the remote LLDP neighbor, the **show lldp neighbors** command output displays the interface name in the **Port ID** field rather than the SNMP index of the interface, which is displayed by default. If you change the default behavior of generating the SNMP index of the interface as the Port ID TLV, you can reenable the default behavior by including the **locally-assigned** statement at the **[edit protocols lldp** port-id-subtype] hierarchy level.

NOTE: Starting with Junos OS Release 12.3, the value of the MIB variable **lldpLocPortId** depends on the SNMP MIB object entity that is used to generate the port ID TLV. If the port ID TLV generation is configured to use the interface name in the **set port-id-subtype interface-name** command, then the value of the MIB variable **lldpLocPortId** is the interface name and not the SNMP index.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, you can configure management interfaces, such as fxp0 or me0, on MX Series routers to send LLDP frames to and receive LLDP frames from neighboring LLDP interfaces.
14.2	Starting with Junos OS Release 14.2, LLDP is supported on extended ports in the Junos Fusion technology.
12.3	Starting with Junos OS Release 12.3R1, you can generate the interface name as the port ID TLV.
12.3	Starting with Junos OS Release 12.3, the value of the MIB variable **lldpLocPortId** depends on the SNMP MIB object entity that is used to generate the port ID TLV.

Related Documentation

- LLDP Overview on page 337
- Tracing LLDP Operations on page 344
- Example: Configuring LLDP on page 342

- *TX Matrix Router Chassis and Interface Names*
- *Monitoring a Routing Matrix with a TX Matrix Router*

Example: Configuring LLDP

The following example configures LLDP on interface **ge-1/1/1** but disables LLDP on all other interfaces, explicitly configures the default values for all automatically enabled features, and configures a value of 30 seconds for the LLDP configuration notification interval and a value of 30 seconds for the physical topology trap interval.

```
[edit]
protocols {
  lldp {
    advertisement-interval 30;
    hold-multiplier 4;
    interface all {
      disable;
    }
    interface ge-1/1/1;
    lldp-configuration-notification-interval 30;
    ptopo-configuration-maximum-hold-time 300;
    ptopo-configuration-trap-interval 30;
    transmit-delay 2;
  }
}
```

You verify operation of LLDP with several show commands:

- **show lldp <detail>**
- **show lldp neighbors** ***interface-name***
- **show lldp statistics** ***interface-name***
- **show lldp local-information**
- **show lldp remote-global-statisitics**

You can clear LLDP neighbor information or statistics globally or on an interface:

- **clear lldp neighbors** ***interface-name***
- **clear lldp statistics** ***interface-name***

You can display basic information about LLDP with the **show lldp detail** command:

```
user@host> show lldp detail
LLDP                    : Enabled
Advertisement interval  : 30 Second(s)
Transmit delay          : 2 Second(s)
Hold timer              : 4 Second(s)
Notification interval   : 30 Second(s)
Config Trap Interval    : 300 Second(s)
Connection Hold timer   : 60 Second(s)


Interface      LLDP       Neighbor count
ge-1/1/1       Enabled    0

LLDP basic TLVs supported:
Chassis identifier, Port identifier, Port description, System name, System
```



```
description, System capabilities, Management address.

LLDP 802 TLVs supported:
Link aggregation, Maximum frame size, MAC/PHY Configuration/Status, Port VLAN ID,
 Port VLAN name.
```

For more details about the output of these commands, see the CLI Explorer.

Related Documentation

- LLDP Overview on page 337
- Configuring LLDP on page 338
- Tracing LLDP Operations on page 344

LLDP Operational Mode Commands

Table 24 on page 343 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot the Link Layer Discovery Protocol (LLDP) protocol. Commands are listed in alphabetical order.

Table 24: LLDP Operational Mode Commands

Task	Command
Clear LLDP neighbor information.	clear lldp neighbors
Clear LLDP statistics.	clear lldp statistics
Display basic LLDP information.	show lldp
Display LLDP local information.	show lldp local-information
Display LLDP neighbor information.	show lldp neighbors
Display LLDP remote global statistics.	show lldp remote-global-statistics
Display LLDP statistics.	show lldp statistics

Related Documentation

- LLDP Overview on page 337
- Configuring LLDP on page 338
- Tracing LLDP Operations on page 344
- Example: Configuring LLDP on page 342

Tracing LLDP Operations

To trace LLDP operational traffic, you can specify options in the global **traceoptions** statement included at the **[edit routing-options]** hierarchy level, and you can specify LLDP-specific options by including the **traceoptions** statement:

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

You can include this statement at the following hierarchy levels:

- **[edit protocols lldp]**
- **[edit routing-instances *routing-instance-name* protocols lldp]**

You can specify the following LLDP-specific options in the LLDP **traceoptions** statement:

- **all**—Trace all operations.
- **config**—Log configuration events.
- **interface**—Trace interface update events.
- **protocol**—Trace protocol information.
- **rtsock**—Trace real-time socket events.
- **vlan**—Trace VLAN update events.

NOTE: Use the trace flag **all** with caution. This flag may cause the CPU to become very busy.

For general information about tracing and global tracing options, see the statement summary for the global **traceoptions** statement in the *Junos OS Routing Protocols Library*.

Related Documentation

- LLDP Overview on page 337
- Configuring LLDP on page 338
- Example: Configuring LLDP on page 342

CHAPTER 15

Configuring VRRP and VRRP for IPv6

VRRP and VRRP for IPv6 Overview

You can configure the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6 for the following interfaces:

- Ethernet
- Fast Ethernet
- Tri-Rate Ethernet copper
- Gigabit Ethernet
- 10-Gigabit Ethernet LAN/WAN PIC
- Ethernet logical interfaces

VRRP and VRRP for IPv6 allow hosts on a LAN to make use of redundant routers on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routers share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routers is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master router, thus always providing a virtual default router and allowing traffic on the LAN to be routed without relying on a single router.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

For VRRP and VRRP for IPv6 overview information, configuration guidelines, and statement summaries, see the *Junos OS High Availability Library for Routing Devices*.

Related Documentation

- Configuring VRRP and VRRP for IPv6 on page 346
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring VRRP and VRRP for IPv6

To configure VRRP or VRRP for IPv6, include the **vrrp-group** or **vrrp-inet6-group** statement, respectively. These statements are available at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

The VRRP and VRRP IPv6 configuration statements are as follows:

```
(vrrp-group | vrrp-inet-group) group-number {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority-number number;
  track {
    priority-hold-time;
    interface interface-name {
      priority-cost priority;
      bandwidth-threshold bits-per-second {
        priority-cost;
      }
    }
  }
  virtual-address [ addresses ];
}
```

You can configure VRRP IPv6 with a global unicast address.

To trace VRRP and VRRP for IPv6 operations, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  file <filename> <files number <match regular-expression <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

When there are multiple VRRP groups, there is a few seconds delay between the time the first gratuitous ARP is sent out and the rest of the gratuitous ARP are sent. Configuring failover-delay compensates for this delay. To configure the failover delay from 500 to 2000 milliseconds for VRRP and VRRP for IPv6 operations, include the **failover-delay *milliseconds*** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
failover-delay milliseconds;
```


To configure the startup period for VRRP and VRRP for IPv6 operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
startup-silent-period seconds;
```

To enable VRRPv3, set the **version-3** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
version-3;
```

Related Documentation

- failover-delay on page 1161
- traceoptions on page 1399
- failover-delay on page 1161
- *vrrp-group*
- VRRP and VRRP for IPv6 Overview on page 345
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 16

Configuring Point-to-Point Protocol over Ethernet

PPPoE Overview

The Point-to-Point Protocol over Ethernet (PPPoE) connects multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

To use PPPoE, you must configure the router as a PPPoE client, encapsulate PPP packets over Ethernet, and initiate a PPPoE session.

M120, M320, and MX Series routers can be configured as a PPPoE access concentrator server. To configure a PPPoE server on an M120, M320, or MX Series Ethernet logical interface, specify PPPoE encapsulation, include the **pp0** statement for the pseudo PPPoE physical interface, and include the **server** statement in the PPPoE options under the logical interface.

NOTE: **PPPoE encapsulation is not supported on M120, M320, or MX Series routers on an ATM2 IQ interface.**

Multiple hosts can be connected to the Services Router, and their data can be authenticated, encrypted, and compressed before the traffic is sent to the PPPoE session on the Services Router's Fast Ethernet or ATM-over-ADSL interface. PPPoE is easy to configure and enables services to be managed on a per-user basis rather than on a per-site basis.

This overview contains the following topics:

PPPoE Interfaces

The PPPoE configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

- If the interface is Fast Ethernet, use a PPPoE encapsulation.
- If the interface is ATM over ADSL, use a PPPoE over ATM encapsulation.

The PPPoE interface on M120 or M320 routers acting as a access concentrator can be a Gigabit Ethernet or 10-Gigabit Ethernet interface.

Ethernet Interface

The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 26 on page 351 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 26: PPPoE Session on an Ethernet Loop

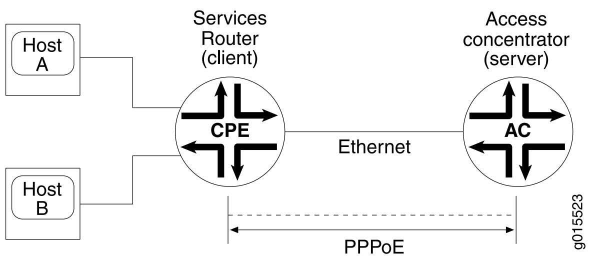

PPPoE Stages

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the discovery stage, the client discovers the access concentrator by identifying the Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the PPPoE session stage, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

NOTE: **If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.**

If you do not configure the access concentrator name, the PPPoE session starts using any available server in the network.

PPPoE Discovery Stage

A Services Router initiates the PPPoE discovery stage by broadcasting a PPPoE active discovery initiation (PADI) packet. To provide a point-to-point connection over Ethernet, each PPPoE session must learn the Ethernet MAC address of the access concentrator and establish a session with a unique session ID. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.

NOTE: **A Services Router cannot receive PPPoE packets from two different access concentrators on the same physical interface.**

The PPPoE discovery stage consists of the following steps:

1. PPPoE active discovery initiation (PADI)—The client initiates a session by broadcasting a PADI packet on the LAN to request a service.
2. PPPoE active discovery offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains it own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.

3. PPPoE active discovery request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.

4. PPPoE active discovery session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session.

   - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.

   - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

PPPoE Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends the PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A Services Router supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions on all interfaces on the Services Router.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE active discovery termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic.

Optional CHAP Authentication

For interfaces with PPPoE encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and be authenticated by its peer.

If you configure an interface to handle incoming CHAP packets only (by including the **passive** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level), the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not include the **passive** statement, the interface always challenges its peer.

For more information about CHAP, see *Configuring the PPP Challenge Handshake Authentication Protocol*.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Evaluation Order for Matching Client Information in PPPoE Service Name Tables*
- *Benefits of Configuring PPPoE Service Name Tables*

- *Configuring PPPoE Service Name Tables*
- *Creating a Service Name Table*
- *Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag*
- *Configuring the Action Taken for the Any Service*
- *Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag*
- *Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information*
- *Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name*
- *Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client*
- *Enabling Advertisement of Named Services in PADO Control Packets*
- *Assigning a Service Name Table to a PPPoE Underlying Interface*
- *Example: Configuring a PPPoE Service Name Table*

- *Troubleshooting PPPoE Service Name Tables*

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring PPPoE

- Configuring the Protocol MTU PPPoE on page 358
- Example: Configuring a PPPoE Server Interface on an M120 or M320 Router on page 359

Overview

To configure PPPoE on an M120 or M320 Multiservice Edge Router or MX Series Universal Edge Router operating as an access concentrator, perform the following tasks:

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
3. Optionally, configure the maximum transmission unit (MTU) of the interface.
4. Configure the operational mode as server.
5. Configure the PPPoE interface address.
6. Configure the destination PPPoE interface address.
7. Optionally, configure the MTU size for the protocol family.
8. Starting in Junos OS Release 10.0, optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.
9. Starting in Junos OS Release 12.3, optionally, disable the sending of PADS messages that contain certain error tags.

NOTE: Starting in Junos OS Release 10.4, when you configure a static PPPoE logical interface, you must include the **pppoe-options** subhierarchy at the **[edit interfaces pp0 unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]** hierarchy level. If you omit the **pppoe-options** subhierarchy from the configuration, the commit operation fails.

Setting the Appropriate Encapsulation on the PPPoE Interface

For PPPoE on an Ethernet interface, you must configure encapsulation on the logical interface and use PPP over Ethernet encapsulation.

For PPPoE on an ATM-over-ADSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL logical interface, use PPPoE over AAL5 LLC encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.

NOTE: PPPoE encapsulation is not supported on an M120 or M320 router on an ATM2 IQ interface.

When you configure a point-to-point encapsulation such as PPP on a physical interface, the physical interface can have only one logical interface (only one **unit** statement) associated with it.

To configure physical interface properties, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation ethernet-over-atm;
```

To configure logical interface encapsulation properties, include the **encapsulation** statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Perform the task appropriate for the interface on which you are using PPPoE. For more information on how to configure PPoE encapsulation on an ethernet interface and on an ATM-over-ADSL interface, see "Configuring PPPoE Encapsulation on an Ethernet Interface" on page 355 and "Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface" on page 355.

Configuring PPPoE Encapsulation on an Ethernet Interface

Both the client and the server must be configured to support PPPoE. To configure PPPoE encapsulation on an Ethernet interface, include the **encapsulation** statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]**

Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

To configure the PPPoE encapsulation on a ATM-over-ADSL interface, perform the following steps:

1. Include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, and specify **ethernet-over-atm**:

   ```
   [edit interfaces pp0]
   encapsulation ethernet-over-atm;
   ```

2. Configure LLC encapsulation on the logical interface by including the **encapsulation** statement and specifying **ppp-over-ether-over-atm-llc**:

   ```
   encapsulation ppp-over-ether-over-atm-llc;
   ```


You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]**

Configuring the PPPoE Underlying Interface

To configure the underlying Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface, include the **underlying-interface** statement:

```
underlying-interface interface-name;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Specify the logical Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface as the underlying interface—for example, **at-0/0/1.0** (ATM VC), **fe-1/0/1.0** (Fast Ethernet interface), or **ge-2/0/0** (Gigabit Ethernet interface).

Identifying the Access Concentrator

When configuring a PPPoE client, identify the access concentrator by a unique name by including the **access-concentrator** statement:

```
access-concentrator name;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Automatic Reconnect Wait Timer

By default, after a PPPoE session is terminated, the session attempts to reconnect immediately. When configuring a PPPoE client, you can specify how many seconds to wait before attempting to reconnect, by including the **auto-reconnect** statement:

```
auto-reconnect seconds;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

You can configure the reconnection attempt to occur in 0 through 4,294,967,295 seconds after the session terminates.

Configuring the PPPoE Service Name

When configuring a PPPoE client, identify the type of service provided by the access concentrator—such as the name of the Internet service provider (ISP), class, or quality of service—by including the **service-name** statement:

```
service-name name;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Server Mode

When configuring a PPPoE server, identify the mode by including the **server** statement:

```
server;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Client Mode

When configuring a PPPoE client, identify the mode by including the **client** statement:

```
client;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Source and Destination Addresses

When configuring a PPPoE client or server, assign source and destination addresses—for example, **192.168.1.1/32** and **192.168.1.2**. To assign the source and destination address, include the **address** and **destination** statements:

```
address address {
  destination address;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces pp0.0 family inet]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family inet]**

Deriving the PPPoE Source Address from a Specified Interface

For a router supporting PPPoE, you can derive the source address from a specified interface—for example, the loopback interface, **lo0.0**—and assign a destination address—for example, **192.168.1.2**. The specified interface must include a logical unit number and have a configured IP address. To derive the source address and assign the destination address, include the **unnumbered-address** and **destination** statements:

```
unnumbered-address interface-name destination address;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces pp0.0 family inet]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family inet]**

Configuring the PPPoE IP Address by Negotiation

You can have the PPPoE client router obtain an IP address by negotiation with the remote end. This method might require the access concentrator to use a RADIUS authentication server. To obtain an IP address from the remote end by negotiation, include the **negotiate-address** statement:

```
negotiate-address;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0.0 family (inet | inet6 | mpls)]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family (inet | inet6 | mpls)]**

Configuring the Protocol MTU PPPoE

You can configure the maximum transmission unit (MTU) size for the protocol family. Specify a range from 0 through 5012 bytes. Ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. To set the MTU, include the **mtu** statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0.0 family (inet | inet6 | mpls)]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family (inet | inet6 | mpls)]**

You can modify the MTU size of the interface by including the **mtu *bytes*** statement at the **[edit interfaces pp0]** hierarchy level:

```
[edit interfaces pp0]
mtu bytes;
```

The default media MTU size used and the range of available sizes on a physical interface depends on the encapsulation used on that interface.

Example: Configuring a PPPoE Server Interface on an M120 or M320 Router

Configure a PPPoE server over a Gigabit Ethernet interface:

```
[edit interfaces]
ge-1/0/0 {
   vlan-tagging;
   unit 1 {
      encapsulation ppp-over-ether;
      vlan-id 10;
   }
}
pp0 {
   unit 0 {
      pppoe-options {
         underlying-interface ge-1/0/0.0;
         server;
      }
      ppp-options {
      }
      family inet {
         address 22.2.2.1/32 {
            destination 22.2.2.2;
         }
      }
   }
}
```

Release History Table

Release	Description
12.3	Starting in Junos OS Release 12.3, optionally, disable the sending of PADS messages that contain certain error tags.
10.4	Starting in Junos OS Release 10.4, when you configure a static PPPoE logical interface, you must include the **pppoe-options** subhierarchy at the **[edit interfaces pp0 unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]** hierarchy level.
10.0	Starting in Junos OS Release 10.0, optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.

Related Documentation

- PPPoE Overview on page 350
- *Ethernet Interfaces Feature Guide for Routing Devices*

Disabling the Sending of PPPoE Keepalive Messages

When configuring the client, you can disable the sending of keepalive messages on a logical interface by including the **no-keepalives** statement:

```
no-keepalives;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]**

Related Documentation

- PPPoE Overview on page 350
- *Ethernet Interfaces Feature Guide for Routing Devices*

Verifying a PPPoE Configuration

Purpose You can use show commands to display and verify the PPPoE configuration.

Action To verify a PPPoE configuration, you can issue the following operational mode commands:

- **show interfaces at-*fpc/pic/port* extensive**
- **show interfaces pp0**
- **show pppoe interfaces**
- **show pppoe version**
- **show pppoe service-name-tables**
- **show pppoe sessions**
- **show pppoe statistics**
- **show pppoe underlying-interfaces**

For more information about these operational mode commands, see CLI Explorer.

Related Documentation

- PPPoE Overview on page 350
- *Ethernet Interfaces Feature Guide for Routing Devices*

Tracing PPPoE Operations

The Junos OS trace feature tracks PPPoE operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file called **pppoed** located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located.

2. When the file **pppoed** reaches 128 kilobytes (KB), it is renamed **pppoed.0**, then **pppoed.1**, and finally **pppoed.2**, until there are three trace files. Then the oldest trace file (**pppoed.2**) is overwritten.

   You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the System Log Explorer.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPPoE tracing operations:

1. Specify that you want to configure tracing options.

   ```
   [edit protocols pppoe]
   user@host# edit traceoptions
   ```

2. (Optional) Configure the name for the file used for the trace output.
3. (Optional) Configure the number and size of the log files.
4. (Optional) Configure access to the log file.
5. (Optional) Configure a regular expression to filter logging events.
6. (Optional) Configure flags to filter the operations to be logged.

Optional PPPoE traceoptions operations are described in the following sections:

- Configuring the PPPoE Trace Log Filename on page 361
- Configuring the Number and Size of PPPoE Log Files on page 362
- Configuring Access to the PPPoE Log File on page 362
- Configuring a Regular Expression for PPPoE Lines to Be Logged on page 362
- Configuring the PPPoE Tracing Flags on page 362

Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is **pppoed**. You can specify a different name with the **file** option.

See Also

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

See Also

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

See Also

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

See Also

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

- [edit protocols pppoe traceoptions]
  user@host# **set flag authentication**

See Also

- Tracing PPPoE Operations on page 360

Related Documentation

- PPPoE Overview on page 350
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is **pppoed**. You can specify a different name with the **file** option.

Related Documentation

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

Related Documentation

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

Related Documentation

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

Related Documentation

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

- [edit protocols pppoe traceoptions]
  user@host# **set flag authentication**

Related Documentation

- Tracing PPPoE Operations on page 360

Configuring the Severity Level to Filter Which PPPoE Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all** or **verbose**. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages . By default, the trace operation output includes only messages with a severity level of **error**.

To configure the type of messages to be logged:

- Configure the message severity level.

  [edit protocols pppoe]

user@host# **set level *severity***

Related Documentation

- Tracing PPPoE Operations on page 360
- **traceoptions (PPPoE) on page 1411**

CHAPTER 17

Configuring Restricted and Unrestricted Proxy ARP

Restricted and Unrestricted Proxy ARP Overview

By default, the Junos OS responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet Interfaces, you can configure the router or switches to proxy-reply to the ARP requests using the restricted or unrestricted proxy ARP configuration.

You might want to configure restricted or unrestricted proxy ARP for routers that act as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.

NOTE: **From Junos OS Release 10.0 onward, Junos OS does not respond to proxy ARP requests with the default route 0.0.0.0. This behavior is in compliance with RFC 1027.**

Restricted Proxy ARP

Restricted proxy ARP enables the router or switch to respond to the ARP requests in which the physical networks of the source and target are not the same and the router or switch has an active route to the target address in the ARP request. The router does not reply if the target address is on the same subnet and the same interface as the ARP requestor.

Unrestricted Proxy ARP

Unrestricted proxy ARP enables the router or switch to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.

WARNING: If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2 reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.

NOTE: While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

Topology Considerations for Unrestricted Proxy ARP

In most situations, you should not configure the router or switch to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used. Figure 27 on page 368 and Figure 28 on page 369 show examples of situations in which you might want to configure unrestricted proxy ARP.

In Figure 27 on page 368, the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In Figure 28 on page 369, the Broadband Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

Figure 27: Edge Device Case for Unrestricted Proxy ARP

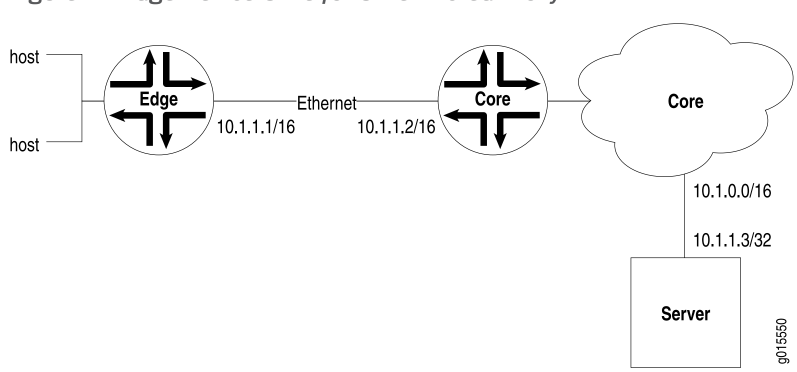

Figure 28: Core Device Case for Unrestricted Proxy ARP

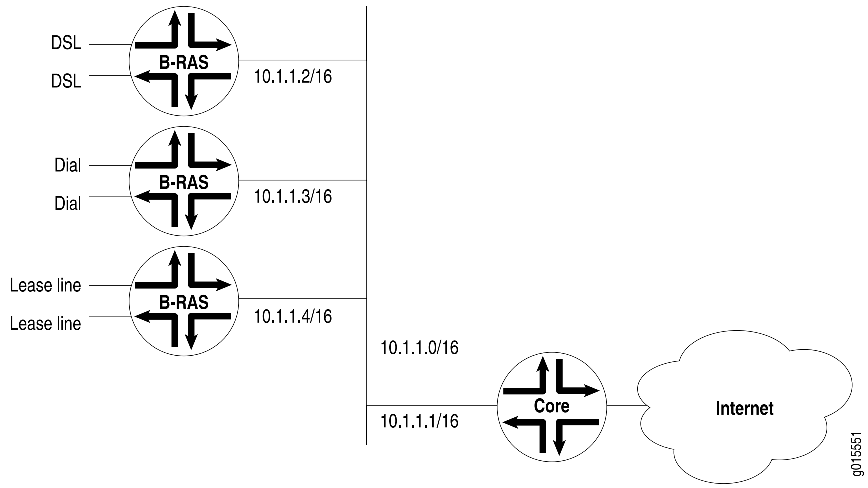

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 369
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

```
proxy-arp (restricted |unrestricted);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.

NOTE: When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See “Configuring Gratuitous ARP” on page 19 for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- proxy-arp on page 1331
- Restricted and Unrestricted Proxy ARP Overview on page 367
- Configuring Gratuitous ARP on page 19
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 18

Configuring Static ARP Table Entries

Static ARP Table Entries Overview

For Fast Ethernet, Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure static ARP table entries, defining mappings between IP and MAC addresses.

Related Documentation

- Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses on page 372

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Also, unlike dynamically learned ARP entries, static ARP entries do not age out. You can also configure static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

To configure static ARP entries:

1. In the configuration mode, at the [**edit**] hierarchy level, configure the router interface on which the ARP table entries for the router is configured.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Configure the protocol family, the logical unit of the interface, and the interface address of the router interface at the [**edit interfaces *interface-name***] hierarchy level. While configuring the protocol family, specify **inet** as the protocol family.

   NOTE: When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

   ```
   [edit interfaces interface-name]
   user@host# edit unit logical-unit-number family inet address interface-address
   ```

3. Configure a static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. The IP address specified must be part of the subnet defined in the enclosing **address** statement. The MAC address must be specified as hexadecimal bytes in the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn*** format. For instance, you can use either **0011.2233.4455** or **00:11:22:33:44:55**.

[edit interfaces *interface-name* unit *logical-unit-number* family inet address
interface-address
user@host# set arp *ip-address* mac *mac-address*

4. Configure another static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. You can also associate a multicast MAC address with a unicast IP address by including the **multicast-mac** option with the **arp** statement. You can optionally configure the router to respond to ARP requests for the specified IP address by using the **publish** option with the **arp** statement.

   NOTE: For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

   [edit interfaces *interface-name* unit *logical-unit-number* family inet address
   interface-address
   user@host# set arp *ip-address* multicast-mac *mac-address* publish

NOTE: The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

Related Documentation

- arp on page 1082
- Static ARP Table Entries Overview on page 371
- Management Ethernet Interface Overview on page 25
- *Interfaces Overview for Switches*
- *Applying Policers*
- *Configuring an Unnumbered Interface*
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 19

Configuring TCC and Layer 2.5 Switching

TCC and Layer 2.5 Switching Overview

Translational cross-connect (TCC) allows you to forward traffic between a variety of Layer 2 protocols or circuits. It is similar to its predecessor, CCC. However, while CCC requires the same Layer 2 encapsulations on both sides of a router (such as Point-to-Point Protocol [PPP] or Frame Relay-to-Frame Relay), TCC lets you connect different types of Layer 2 protocols interchangeably. With TCC, combinations such as PPP-to-ATM and Ethernet-to-Frame Relay cross-connections are possible. Also, TCC can be used to create Layer 2.5 VPNs and Layer 2.5 circuits.

Consider a sample topology (Figure 29 on page 375) in which you can configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. In this topology, Router B strips all PPP encapsulation data from frames arriving from Router A and adds ATM encapsulation data before the frames are sent to Router C. All Layer 2 negotiations are terminated at the interconnecting router (Router B).

Figure 29: Sample Translation Cross-Connect Topology

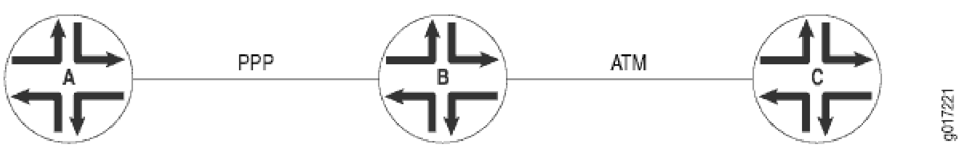

TCC functionality is different from standard Layer 2 switching. TCC only swaps Layer 2 headers. No other processing, such as header checksums, time-to-live (TTL) decrementing, or protocol handling, is performed. Currently, TCC is supported in IPv4, ISO, and MPLS.

Ethernet TCC is supported on interfaces that carry IPv4 traffic only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC and extended VLAN CCC are not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.

Related Documentation

- Configuring VLAN TCC Encapsulation on page 376
- Configuring Translation Cross-Connect Interface Switching on page 378

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring VLAN TCC Encapsulation

VLAN TCC encapsulation allows circuits to have different media on either side of the forwarding path. VLAN TCC encapsulation supports TPID 0x8100 only. You must include configuration statements at the logical and physical interface hierarchy levels.

To configure VLAN TCC encapsulation, include the **encapsulation** statement and specify the **vlan-tcc** option:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-tcc;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Additionally, configure the logical interface by including the **proxy** and **remote** statements:

```
proxy {
  inet-address;
}
remote {
  (inet-address | mac-address);
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family tcc]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]**

The proxy address is the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy.

The remote address is the IP or MAC address of the remote router. The **remote** statement provides ARP capability from the TCC switching router to the Ethernet neighbor. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

When VLAN TCC encapsulation is configured on the logical interface, you also must specify flexible Ethernet services on the physical interface. To specify flexible Ethernet services, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level and specify the **flexible-ethernet-services** option:

```
[edit interfaces interface-name]
encapsulation flexible-ethernet-services;
```

Extended VLAN TCC encapsulation supports TPIDs 0x8100 and 0x9901. Extended VLAN TCC is specified at the physical interface level. When configured, all units on that interface must use VLAN TCC encapsulation, and no explicit configuration is needed on logical interfaces.

One-port Gigabit Ethernet, 2-port Gigabit Ethernet, and 4-port Fast Ethernet PICs with VLAN tagging enabled can use VLAN TCC encapsulation. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level and specify the **extended-vlan-tcc** option:

```
[edit interfaces interface-name]
encapsulation extended-vlan-tcc;
```

For VLAN TCC encapsulation, all VLAN IDs from 1 through 1024 are valid. VLAN ID 0 is reserved for tagging the priority of frames.

Extended VLAN TCC is not supported on 4-port Gigabit Ethernet PICs.

Related Documentation

- **encapsulation on page 1135**
- **remote on page 1340**
- **proxy on page 1330**
- TCC and Layer 2.5 Switching Overview on page 375
- Configuring Translation Cross-Connect Interface Switching on page 378
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Translation Cross-Connect Interface Switching

To configure a full-duplex Layer 2.5 translation cross-connect between two routers (A and C), you can configure a Juniper Networks router (Router B) as the TCC interface. Ethernet TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. Consider the topology in Figure 30 on page 378 where the Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit accepts packets carrying standard TPID values.

Figure 30: Sample Topology of Layer 2.5 Translational Cross-Connect

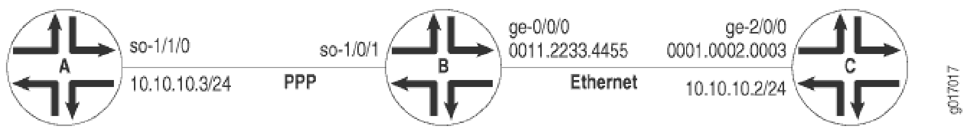

If traffic flows from Router A to Router C, the Junos OS strips all PPP encapsulation data from incoming packets and adds Ethernet encapsulation data before forwarding the packets. If traffic flows from Router C to Router A, the Junos OS strips all Ethernet encapsulation data from incoming packets and adds PPP encapsulation data before forwarding the packets.

To configure the router as the translational cross-connect interface:

1. In the configuration mode, at the [**edit**] hierarchy level, first configure the interface that is connected to Router A.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. (Optional) Specify the description of the interface. For example, you could specify the interface name on Router A that is connected to this interface.

   ```
   [edit interfaces interface-name]
   user@host# set description description
   ```

3. Specify the encapsulation. If the Router A to Router B circuit is PPP, then specify **ppp-tcc** as the encapsulation. If the Router A to Router B circuit is frame relay, specify **frame-relay-tcc**.

   ```
   [edit interfaces interface-name]
   user@host# set encapsulation encapsulation-type
   ```

4. In the configuration mode, at the [**edit**] hierarchy level, first configure the interface that is connected to Router C.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

5. (Optional) Specify the description of this interface. For example, you could specify the interface name on Router C that is connected to this interface.

```
[edit interfaces interface-name]
user@host# set description description
```

6. Specify the encapsulation. If the Router B to Router C circuit is Ethernet, then specify **ethernet-tcc** as the encapsulation. If the Router B to Router C circuit is ATM, specify **atm-tcc-vc-mux**.

```
[edit interfaces interface-name]
user@host# set encapsulation encapsulation-type
```

7. Specify the IP address or MAC address of the remote router to provide address resolution protocol (ARP) for the TCC router's Ethernet-based neighbor using the **remote** statement. You must specify the statement at the [**edit interfaces interface-name unit unit-number family tcc**] hierarchy level. You can a specify the MAC address of the remote router instead of the IP address. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

```
[edit interfaces interface-name]
user@host# set unit 0 family family remote inet-address ip-address
```

8. Specify the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy using the **proxy** statement. You must specify the statement at the [**edit interfaces interface-name unit unit-number family tcc**] hierarchy level.

```
[edit interfaces interface-name]
user@host# set unit 0 family family proxy inet-address ip-address
```

To verify the TCC connection, use the **show connections** command on TCC router.

Related Documentation

- **encapsulation on page 1135**
- **remote on page 1340**
- **proxy on page 1330**
- TCC and Layer 2.5 Switching Overview on page 375
- Configuring VLAN TCC Encapsulation on page 376

- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 20

Configuring Link Degrade Monitoring

Link Degrade Monitoring Overview

Link degrade monitoring enables you to monitor the quality of physical links on Ethernet interfaces (10-Gigabit, 40-Gigabit, and 100-Gigabit) and take corrective action when the link quality degrades beyond a certain level. You can configure this feature by applying the **link-degrade-monitor** statement at the **[edit interfaces *interface-name*]** hierarchy level. When configured on your device, this feature continuously monitors bit error rate (BER) value of the link and initiates a corrective action (media-based) when the BER value breaches a user-configured threshold. The feature can detect a BER value as low as 10^{-13} through 10^{-5}, helping you prevent or minimize packet drops in physical links.

You can configure autorecovery or manual recovery method for the degraded link. In the case of manual recovery, you need to use the **request interface link-degrade-recover** *interface-name* statement to recover the degraded link. If autorecovery is configured, automatic recovery of the degraded link is attempted at the user configured intervals, and when the link's BER value is within the configured limit, the link is recovered.

NOTE: Layer 2 and Layer 3 protocols already support physical link monitoring. So do Ethernet links through the Link Fault Signaling (LFS) protocol. However, these existing mechanisms cannot detect BER ranges that are very low (for example, 10^{-13} through 10^{-5}).

Supported Platforms

Table 25 on page 382 lists the platform series and line cards that support link degrade monitoring.

Table 25: Line Cards that Support Link Degrade Monitoring

Platform Series	MPC Line Cards Supported	DPC Line Cards Supported

MX	• MPC4E-3D-2CGE-8XGE
	• MPC4E-3D-32XGE-SFPP
	• MPC-3D-16XGE-SFP
	• MPC3 with MIC3-3D-1X100GE-CFP
	• MPC3 with MIC3-3D-2X40GE-QSFPP
	• MPC3 with MIC-3D-2XGE-XFP
	• MPC3 with 2x10GE XFP MIC
	• MPC3 with 2x10GE XFP MIC
	• MPC5 with following variants:
	  ◦ 2CGE + 4XGE
	  ◦ 24XGE + 6XLGE
	• MPC6 with the following variants:
	  ◦ 2X100GE CFP2
	  ◦ 24X10GE SFPP
	  ◦ 24X10GE SFPP OTN
	  ◦ 4x100GE CXP
		• DPCE-R-Q-4XGE-XFP
		• DPCE-R-4XGE-XFP
		• DPCE-X-4XGE-XFP
		• DPCE-X-Q-4XGE-XFP
		• DPCE-R-2XGE-XFP
		• DPCE-R-4XGE-XFP
		• DPCE-X-4XGE-XFP
		On 10-Gigabit Ethernet interfaces:
		• DPCE-R-Q-20GE-2XGE
		• DPCE-R-20GE-2XGE
		• DPCE-X-20GE-2XGE

Related Documentation

- *Physical Interface Damping Overview*
- *Fast Reroute Overview*
- link-degrade-monitor on page 1221
- thresholds on page 1395
- recovery on page 1339
- request interface link-degrade-recover on page 1495

CHAPTER 21

Configuring Power-over-Ethernet on ACX Series

Understanding PoE on ACX Series Universal Metro Routers

Power over Ethernet (PoE) is the implementation of the IEEE 802.3af and IEEE 802.3at standards that allows both data and electrical power to pass over a copper Ethernet LAN cable.

Juniper Networks provides PoE on ACX2000 Universal Metro Routers that allows power delivery up to 65 W per PoE port. PoE ports transfer electrical power and data to remote devices over standard twisted-pair cables in an Ethernet network. Using the PoE ports, you can plug in devices that require both network connectivity and electrical power, such as voice over IP (VoIP) and wireless LAN access points.

You can configure the ACX2000 Universal Metro Router to act as a power sourcing equipment (PSE), supplying power to powered devices that are connected on designated ports.

This topic contains the following sections: :

ACX2000 PoE Specifications

Table 26 on page 384 lists the PoE specifications for the ACX2000 routers.

Table 26: PoE Specifications for the ACX2000 Routers

Specifications	For ACX2000 Universal Metro Routers
Supported standards	• IEEE 802.3 AF • IEEE 802.3 AT (PoE+) • Legacy (pre-standards)
Supported ports	Supported on only two Gigabit Ethernet ports (ge-0/1/3 and ge-0/1/7).
Total PoE power sourcing capacity	130 W
Default per port power limit	32 W
Maximum per port power limit	65 W
Power management modes	• **class**—Power allocated for each interface can be configured. • **static**—Power allocated for interfaces is based on the class of powered device connected. • **high-power**—Power allocated for interfaces up to 65 W per port.

PoE Classes and Power Ratings

A powered device is classified based on the maximum power that it draws across all input voltages and operational modes. When class-based power management mode is configured on the ACX2000 routers, power is allocated taking into account the maximum power ratings defined for the different classes of devices.

Table 27 on page 384 lists the classes and their power ratings as specified by the IEEE standards.

Table 27: ACX2000 Universal Metro Router PoE Specifications

Class	Usage	Minimum Power Levels Output from PoE Port
0	Default	15.4 W
1	Optional	4.0 W
2	Optional	7.0 W
3	Optional	15.4 W
4	Reserved	Class 4 power devices are eligible to receive power up to 30 W according to the IEEE standards.

PoE Options

For ACX2000 Universal Metro Routers that support PoE ports, the factory default configuration enables PoE on the PoE-capable ports, with default settings in effect. You might not have to do any additional configuration if the default settings work for you. Table 28 on page 385 shows the PoE configuration options and their default settings for the PoE controller and for the PoE interfaces.

Table 28: PoE Configuration Options and Default Settings

Option	Default	Description
PoE Controller Options		
guard-band	0 W	Reserves up to 19 W power from the PoE power budget to be used in the case of a spike in PoE power consumption.
management	static	Sets the PoE power management mode for the router. The power management mode determines how power to a PoE interface is allocated: • **class**—Power allocated for each interface can be configured. • **static**—Power allocated for interfaces is based on the class of powered device connected. • **high-power**—Power allocated for interfaces up to 65 W per port.
Interface Options		
disable (Power over Ethernet)	Not included in default configuration	When included in the configuration, disables PoE on the interface. The interface maintains network connectivity but no longer supplies power to a connected powered device. Power is not allocated to the interface.
priority (Power over Ethernet)	low	Sets an interface's power priority to either **low** or **high**. If power is insufficient for all PoE interfaces, the PoE power to low-priority interfaces is shut down before power to high-priority interfaces is shut down. Among interfaces that have the same assigned priority, the power priority is determined by port number, with lower-numbered ports having higher priority.
telemetries	Not included in default configuration	When included in the configuration, enables the logging of power consumption records on an interface. Logging occurs every 5 minutes for 1 hour unless you specify a different value for ***interval (Power over Ethernet)*** or ***duration***.

Related Documentation

- Example: Configuring PoE on ACX2000 Routers on page 385
- Example: Disabling a PoE Interface on ACX2000 Routers on page 390

Example: Configuring PoE on ACX2000 Routers

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that need both network connectivity and electric power, such as voice over IP (VoIP) phones, wireless access points, and IP cameras.

This example shows how to configure PoE to deliver power up to 65 W on ACX2000 interfaces:

- Requirements on page 386
- Overview on page 386
- Configuration on page 387
- Verification on page 388

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2 or later for ACX Series routers
- An ACX2000 router that supports PoE

Before you configure PoE, be sure you have:

- Performed the initial router configuration. See *ACX Series Autoinstallation Overview*,*Verifying Autoinstallation on ACX Series Universal Metro Routers*, and *Boot Sequence on ACX Series Routers* for details.

Overview

This example consists of a router that has eight ports. Only two ports—ge-0/1/3 and ge-0/1/7—support PoE, which means they provide both network connectivity and electric power for powered devices such as VoIP telephones, wireless access points, and IP security cameras that require power up to 65 W. The remaining six ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers.

Table 29 on page 386 details the topology used in this configuration example.

Table 29: Components of the PoE Configuration

Property	Settings
Hardware	ACX2000 router with 8 Gigabit Ethernet ports: Two PoE interfaces (ge-0/1/3 and ge-0/1/7) and 6 non-PoE interfaces (ge-0/1/0, ge-0/1/1, ge-0/1/2, ge-0/1/4, ge-0/1/5, ge-0/1/6).
VLAN name	default
Connection to a wireless access point (requires PoE)	ge-0/1/7
Power port priority	high
Maximum power available to PoE port	65 W
PoE management mode	high-power

Table 29: Components of the PoE Configuration (continued)

Property	Settings
Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)	ge-0/1/0 through ge-0/1/2
Unused ports (for future expansion)	ge-0/1/4 through ge-0/1/6

Configuration

To configure PoE on an ACX2000 router:

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set poe management high-power guard-band 19
set poe interface ge-0/1/3 priority high maximum-power 65 telemetries
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PoE:

1. Set the PoE management mode to **high-power**.

   ```
   [edit]
   user@host# set poe management high-power
   ```

   NOTE:
   - Set the PoE management mode to **high-power** only when the power requirement is more than 32 W and up to 65 W. If the power requirement is less than or equal to 32 W, then you do not need to set the PoE management mode to **high-power**.
   - The default management mode is **static**. In this mode, the power sourcing equipment can deliver power up to 32 W.

2. Reserve power wattage in case of a spike in PoE consumption.

   ```
   [edit]
   user@host# set poe guard-band 19
   ```

3. Enable PoE.

   ```
   [edit]
   user@host# edit poe interface ge-0/1/3
   ```

4. Set the power port priority.

```
[edit poe interface ge-0/1/3]
user@host# set priority high
```

5. Set the maximum PoE power for a port.

```
[edit poe interface ge-0/1/3]
user@host# set maximum-power 65
```

NOTE: Set the maximum PoE power for a port only when the power requirement is more than 32 W and up to 65 W. If the power requirement is less than or equal to 32 W, then you do not need to configure the maximum PoE power.

6. Enable the logging of PoE power consumption.

```
[edit poe interface ge-0/1/3]
user@host# set telemetries
```

Results

In configuration mode, confirm your configuration by entering the **show poe interface ge-0/1/3** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show poe interface ge-0/1/3
  priority high;
  maximum-power 65;
telemetries;
```

If you are done configuring the device, enter **commit** in configuration mode.

Verification

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

- Verifying the Status of PoE Interfaces on page 388
- Verifying the Telemetry Data (History) for the Specified Interface on page 389
- Verifying PoE Global Parameters on page 389

Verifying the Status of PoE Interfaces

Purpose Verify that the PoE interfaces are enabled and set to the desired priority settings.

Action In operational mode, enter the **show poe interface ge-0/1/3** command.

```
user@host> show poe interface ge-0/1/3
PoE interface status:
PoE interface                : ge-0/1/3
Administrative status        : Enabled
Operational status           : Powered-up
Power limit on the interface : 65 W
Priority                     : High
Power consumed               : 6.6 W
Class of power device        : 0
```

Meaning The **show poe interface ge-0/1/3** command lists PoE interfaces configured on the ACX2000 router, with their status, priority, power consumption, and class.

Verifying the Telemetry Data (History) for the Specified Interface

Purpose Verify the PoE interface's power consumption over a specified period.

Action In operational mode, enter the **show poe telemetries interface** command.

For all records:

```
user@host> show poe telemetries interface ge-0/1/3 all
Interface  Sl No    Timestamp                Power    Voltage
              1     Mon May 14 00:45:05 2012 14.2 W   53.9 V
              2     Mon May 14 00:44:04 2012 14.2 W   53.9 V
              3     Mon May 14 00:43:03 2012 14.2 W   53.9 V
```

For a specific number of records:

```
user@host> show poe telemetries interface ge-0/1/3 2
Interface  Sl No    Timestamp                Power    Voltage
              1     Mon May 14 00:45:05 2012 14.2 W   53.9 V
              2     Mon May 14 00:44:04 2012 14.2 W   53.9 V
```

Meaning The telemetry status displays the power consumption history for the specified interface, provided telemetry has been configured for that interface.

Verifying PoE Global Parameters

Purpose Verify global parameters such as guard band, power limit, and power consumption.

Action In operational mode, enter the **show poe controller** command.

```
user@host> show poe controller
Controller  Maximum  Power        Guard    Management   Status        Lldp
index       power    consumption  band                               Priority
```

```
0         130.0 W   14.2 W          0 W      high-power    UP
```

Meaning The **show poe controller** command lists the global parameters configured on the router.

Related Documentation

- Understanding PoE on ACX Series Universal Metro Routers on page 383

Example: Disabling a PoE Interface on ACX2000 Routers

This example shows how to disable PoE on all interfaces or on a specific interface.

Requirements

Before you begin:

- Configure PoE on all interfaces. See "Example: Configuring PoE on ACX2000 Routers" on page 385.

Overview

In this example, you disable PoE on all interfaces and on a specific interface, which in this case is ge-0/1/3.

Configuration

Step-by-Step Procedure

- Disable PoE on all interfaces.

  [edit]
  user@host# **set poe interface all disable**

- Disable PoE on a specific interface.

  [edit]
  user@host# **set poe interface ge-0/1/3 disable**

Verification

To verify the configuration is working properly, enter the **show poe interface** command.

```
user@host> show poe interface

Interface     Admin       Oper     Max        Priority       Power          Class
              status      status   power                     consumption
 ge-0/1/3     Disabled    Disabled 32.0W      Low            0.0W            0
 ge-0/1/7     Disabled    Disabled 32.0W      Low            0.0W            0
```



```
user@host> show poe interface ge-0/1/3

PoE interface status:
PoE interface                : ge-0/1/3
Administrative status        : Disabled
Operational status           : Disabled
Power limit on the interface : 32.0 W
Priority                     : Low
Power consumed               : 0.0 W
Class of power device        : 0
```

Related Documentation

- Understanding PoE on ACX Series Universal Metro Routers on page 383

Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers

Problem

Description: A Power over Ethernet (PoE) interface is not supplying power to the powered device.

Solution

Check for the items shown in Table 30 on page 391.

Table 30: Troubleshooting a PoE Interface

Items to Check	Explanation
Is interface PoE enabled?	Only interfaces **ge-0/1/3** and **ge-0/1/7** can function as PoE ports.
Has PoE capability been disabled for that interface?	Use the **show poe interface** command to check PoE interface status.
Is the cable properly seated in the port socket?	Check the hardware.
Does the powered device require more power than is available on the interface?	Use the **show poe interface** command to check the maximum power provided by the interface.
If the ***telemetries*** option has been enabled for the interface, check the history of power consumption.	Use the **show poe telemetries** command to display the history of power consumption.

Related Documentation

- Understanding PoE on ACX Series Universal Metro Routers on page 383
- Example: Configuring PoE on ACX2000 Routers on page 385

PART 2

Gigabit Ethernet Interfaces

CHAPTER 22

Configuring 10-Gigabit Ethernet LAN/WAN PICs

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview

This section describes the main features and caveats of the 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP) and specifies which routers support this PIC.

The 10–port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) is supported on Juniper Networks T640 Core Routers, T1600 Core Routers, and T4000 Core Routers. It has the following features:

- Access to all 10-Gigabit Ethernet port counters through SNMP
- Intelligent handling of oversubscribed traffic in applications such as data centers and dense-core uplinks
- Line-rate operation for five 10-Gigabit Ethernet ports from each port group, or a total WAN bandwidth of 100 Gbps with Packet Forwarding Engine bandwidth of 50 Gbps
- Flexible encapsulation, source address and destination address media access control (MAC) filtering, source address MAC learning, MAC accounting, and MAC policing

- Interface encapsulations, such as the following:
  - **ethernet-ccc**—Ethernet cross-connect
  - **vlan-ccc**—802.1Q tagging for a cross-connect
  - **ethernet-tcc**—Ethernet translational cross-connect
  - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect
  - **extended-vlan-ccc**—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect
  - **ethernet-vpls**—Ethernet virtual private LAN service
  - **vlan-vpls**—VLAN virtual private LAN service
  - **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration
- WAN PHY features, such as the following:
  - WAN PHY mode on a per-port basis
  - Insertion and detection of path trace messages
  - Ethernet WAN Interface Sublayer (WIS) object

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).

- Single, stacked, and flexible VLAN tagging modes
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames
- Link aggregation group (LAG) on single chassis
- Interoperability with other 10-Gigabit Ethernet PICs in M Series and T Series routers in the LAN PHY and WAN PHY modes
- Interrupt-driven link-down detection mechanism
- Two-to-one oversubscription of traffic across a port group

  Traffic from 10 ingress ports to the Packet Forwarding Engine traffic is statically mapped to one of the 5 egress ports. 10 Gbps of bandwidth toward the Packet Forwarding Engine is shared by two ingress ports (called a *port group*), thereby achieving two-to-one oversubscription. This scheme provides two-to-one oversubscription across a port group and not across the entire PIC.

- Four queues per physical interface on ingress and eight queues per physical interface on egress
- A separate control queue per physical interface to ensure that the control packets are not dropped during oversubscribed traffic. The control queue can be disabled in the CLI.
- Optical diagnostics
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification
- Weighted round-robin scheduling with two queue priorities (low and strict-high)
- Committed information rate and peak information rate shaping on a per-queue basis
- Excess information rate configuration for allocation of excess bandwidth
- IEEE 802.3ah Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
  - Link fault management
  - Link discovery
  - Graceful Routing Engine Switchover
- IEEE 802.3ag Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
  - Connectivity fault management (CFM)
  - Linktrace
  - Loopback
  - Graceful Routing Engine switchover (GRES)

The 10-port 10-Gigabit Ethernet LAN/WAN PIC has the following caveats:

- Source address and destination address MAC filtering takes place after oversubscription is handled.
- Oversubscription on the PIC operates across a port group of two ports and not at the PIC level.
- Queuing is not supported at the logical interface level.
- Committed information rate and peak information rate configurations are not supported at the physical interface level.
- There is limited packet buffering of 2 MB.
- Delay-bandwidth buffering configuration is not supported.
- Multifield classifiers are not supported at the PIC level.

  The multifield classification can be done at the Packet Forwarding Engine using the firewall filters, which overrides the classification done at the PIC level. The multifield classification at the Packet Forwarding Engine occurs after the PIC handles the oversubscribed traffic.

- Egress MAC policer statistics not supported.
- Byte counters are not supported at the queue level.
- Only TPID (0x8100) is supported.
- Line-timing mode is not supported.
- MAC-level Rx VLAN tagged frames counter is not supported.
- OAM unified in-service software upgrade (unified ISSU) is not supported.
- OAM remote loopback is not supported.

The 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) supports link aggregation. For bandwidth aggregation, load sharing, and link protection, LAG can be enabled. Once aggregated Ethernet is enabled, Link Aggregation Control Protocol (LACP) forms an aggregated bundle of member links.

Only features that are supported across all of the linked devices will be supported in the resulting LAG bundle. The following caveats apply to LAG bundles that involve 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) ports:

- Non-standard TPID for VLAN tagging is not supported, except for 0x8100.
- The number of user created IFLs is limited to 4065/PIC and 1022/port.
- Classifier tables are limited to 8 for each BA classifier type.
- Forwarding classes are limited to 8.
- The **guaranteed-rate** and **shaping-rate** statements are not supported at the IFD level.
- The **per-unit-scheduler** and **hierarchical-scheduler** statements are not supported.
- Only the **strict-high** and **low** levels of scheduling priorities are supported.
- The **excess-priority** configuration is not supported.
- The **buffer-size** configuration under **schedulers** is not supported.
- WRED is not supported.
- srTCM and trTCM are not supported.
- Shared scheduler mode is not supported.

Table 31 on page 398 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

Table 31: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs

Capability	Support
Maximum VLANs per PIC	4065
Maximum VLANs per port	1022

Table 31: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs (continued)

Capability	Support
MAC learning per port	960
MAC accounting per port	960
MAC filtering per port	960 (64 filters per physical or logical interface) 960 filters across multiple logical interfaces
MAC policers	128 ingress Mac policers 128 egress Mac policers
Classifers	Eight classifiers per PIC for each BA classifier type

Release History Table

Release	Description
12.2	Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
12.1R2	Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP).

Related Documentation

- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 404
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 404
- Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC on page 407
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686

- *Ethernet Interfaces Feature Guide for Routing Devices*

12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC is a 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number, PF-12XGE-SFPP) on T4000 Core Routers.

The following features are supported on the 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC:

- Access to all 10-Gigabit Ethernet port counters through SNMP.
- Logical interface–level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Flexible encapsulation.
- Single, stacked, and flexible VLAN tagging modes.
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces.
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames.
- Link aggregation group (LAG) on single chassis.
- Interoperability with other 10-Gigabit Ethernet PICs on M Series and T Series routers in LAN PHY mode.
- Eight queues per physical interface on egress.
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification.
- Defining the VLAN rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.

  NOTE: **Only the Tag Protocol Identifier (TPID) 0x8100 is supported.**

- Interface encapsulations, such as the following:
  - **untagged**—Default encapsulation, when other encapsulation is not configured.
    - You can configure only one logical interface (unit 0) on the port.
    - You cannot include the **vlan-id** statement in the configuration of the logical interface.
  - **vlan-tagging**—Enable VLAN tagging for all logical interfaces on the physical interface.
  - **stacked-vlan-tagging**—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
  - **ethernet-ccc**—Ethernet cross-connect.
  - **ethernet-tcc**—Ethernet translational cross-connect.
  - **vlan-ccc**—802.1Q tagging for a cross-connect.
  - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect.
  - **extended-vlan-ccc**—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect.
  - **extended-vlan-tcc**—Standard Tag Protocol Identifier (TPID) tagging for an Ethernet translational cross-connect.

- **ethernet-vpls**—Ethernet virtual private LAN service.
- **vlan-vpls**—VLAN virtual private LAN service.
- **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration.

- The following Layer 3 protocols are also supported:
  - IPv4
  - IPv6
  - MPLS
- WAN PHY features, such as the following:
  - WAN PHY mode on a per-port basis.
  - Insertion and detection of path trace messages.
  - Ethernet WAN Interface Sublayer (WIS) object.

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on T4000 routers with 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+.

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC does not support:

- MAC filtering, accounting, and policing for destination MAC at the logical interface level.

NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.
- Multiple TPIDs

Capability	Support
Maximum logical interfaces per PIC	32,000
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.
Classifers	Eight classifiers per PIC for each BA classifier type

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 395

24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

This section describes the main features and caveats of the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP).

The following major software features are supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP):

- Twenty-four 10-Gigabit Ethernet interfaces in two-to-one oversubscription of traffic in oversubscribed mode or 12 ports in line-rate mode. For more information about oversubscribed mode and line-rate mode, see the "Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription" on page 404.
- Traffic is classified as control traffic or best-effort traffic with non-class-of-service-aware tail drops of best-effort traffic in oversubscribed mode.

  The aggregate bandwidth of all the ports together is 120 Gbps. No hard partitioning of bandwidth is done—that is, if one port group is active, it can support 120 Gbps traffic. The bandwidth for best-effort traffic is shared among all the 24 ports.

  Note that the preclassification is restricted to two traffic classes, and is not user-configurable.
- All Junos OS configuration commands supported on the existing 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The output of the **show interfaces extensive** operational mode command now displays preclassification queue counters.
- Line-rate mode operation of the first 12 ports can be achieved by using the **[set chassis fpc *fpc-number* pic *pic-number* linerate-mode]** command. By default, the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ works in oversubscribed mode.
- LAN PHY mode and WAN PHY mode on a per-port basis. WAN PHY mode can be achieved by using the **[set interfaces *interface-name* framing wan-phy]** command.
- WAN PHY features, such as the following:
  - Insertion and detection of path trace messages.
  - Ethernet WAN Interface Sublayer (WIS) object.
- Aggregated Ethernet is supported only in line-rate mode.
- Link aggregation group (LAG) is supported only in line-rate mode.
- 4000 logical interfaces per physical interface and 32,000 logical interfaces per chassis.
- Access to all 10-Gigabit Ethernet port counters through SNMP.

NOTE: Graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are now supported on T4000 routers.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- 12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 399
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 404

Modes of Operation of 10-Gigabit Ethernet PICs

10-Gigabit Ethernet PICs operate in the following modes:

- Line-rate mode—By default, the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP) operates in line-rate mode.

  In a 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP), 12 ports (ports 0–11)can operate in line-rate mode . To configure the PF-24XGE-SFPP PIC to operate in line-rate mode, include the **linerate-mode** statement at the **[edit chassis set fpc *fpc-number* pic *pic-number*]** hierarchy level.

- Oversubscribed mode—In this mode, all ports on the PIC are enabled with two-to-one oversubscription. In a PF-24XGE-SFPP PIC, by default, two-to-one oversubscription of traffic is achieved in oversubscribed mode—Traffic from 24 ingress ports to the Packet Forwarding Engine is statically mapped to one of the 12 egress ports. 10 Gbps of bandwidth traffic moving toward the Packet Forwarding Engine is shared by two ingress ports (called a port group), thereby achieving two-to-one oversubscription. This scheme provides two-to-one oversubscription across a port group and not across the entire PIC.

  NOTE: PF-12XGE-SFPP PIC always operates at line rate.

- Mixed-rate mode or dual-rate mode—Dual-rate mode or mixed-rate mode for PF-24XGE-SFPP allows you to configure a mix of port speeds of 1 Gbps and 10 Gbps. However, on PF-12XGE-SFPP, note that you can configure port speeds of either 1 Gbps and 10 Gbps when the PIC is in line rate mode. You can enable mixed-rate mode and set port speeds with the **mixed-rate-mode** and **speed 1G |10G** statements respectively at the **[edit chassis fpc *x* pic *y*]** hierarchy level. You can disable mixed-rate mode with the **delete chassis fpc *x* pic *y* mixed-rate-mode** statement.

  NOTE: To change the port speed from 10 Gbps to 1 Gbps on the PF-24XGE-SFPP and PF-12XGE-SFPP PICs, SFP optics is required.

Related Documentation

- Configuring Mixed-Rate Mode Operation on page 408
- **mixed-rate-mode on page 1261**

Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription

For 10-Gigabit Ethernet LAN/WAN PICs supporting oversubscription, oversubscribed Ethernet mode is set by default. To configure these PICs in line-rate mode, include the **linerate-mode** statement at the **[edit chassis set fpc *fpc-number* pic *pic-number*]** hierarchy level:

```
[edit chassis]
set fpc fpc-number pic pic-number linerate-mode;
```

To return to the default oversubscribed Ethernet mode, delete the **linerate-mode** statement at the **[edit chassis fpc *fpc-number* pic *pic-number*]** hierarchy level.

NOTE: When the mode of operation of a PIC is changed, the PIC is taken offline and then brought back online immediately.

The following 10-Gigabit Ethernet LAN/WAN PICs support line-rate mode:

- 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP)
- 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP)

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 395
- 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 402
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 404
- Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC

On a 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), a control queue is used to queue all control packets received on an ingress port. This ensures that control protocol packets do not get dropped randomly when there is congestion due to oversubscription. The following control protocols are supported:

- OSPF
- OSPF3

- VRRP
- IGMP
- RSVP
- PIM
- BGP
- BFD
- LDP
- IS-IS
- RIP
- RIPV6
- LACP
- ARP
- IPv6 NDP
- Connectivity fault management (CFM)
- Link fault management (LFM)

These control packets can either terminate locally or transit through the router. The control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. Hence, if transit control traffic is taking too much bandwidth, then it can cause drops on locally terminating control traffic, as shown in Figure 31 on page 405.

Figure 31: Control Queue Rate Limiter Scenario

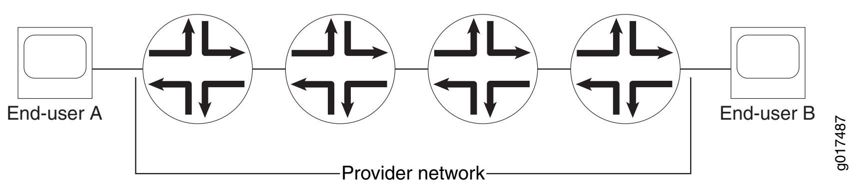

If the end users generate a mass of malicious traffic for which the port number is 179 (BGP), the router dispatches that traffic to the ingress control queue. Further, if congestion occurs in this ingress control queue due to this malicious traffic, the provider's network control packets may be affected.

In some applications, this can be perceived as a new vulnerability. To address this concern, you can disable the control queue feature. With the control queue feature disabled, you must take precautions to protect control traffic through other means, such as mapping control packets (using BA classification) to a queue that is marked strict-high or is configured with a high CIR.

You can disable the control queue for all ports on the PIC. To disable the control queue, use the **set chassis fpc *n* pic *n* no-pre-classifier** command. By default, the **no-pre-classifier** statement is not configured and the control queue is operational.

Deleting the **no-pre-classifier** statement re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

NOTE:

- This functionality is applicable both in OSE and line-rate modes.
- The control queue feature is enabled by default in both OSE and line-rate modes, which can be overridden by the user configuration.
- When the control queue is disabled, various **show queue** commands will show *control queue* in the output. However, all control queue counters are reported as zeros.
- Changing this configuration (enabling or disabling the control queue feature) results in the PIC being taken offline and brought back online.

Once the control queue is disabled, the Layer 2/Layer 3 control packets are subject to queue selection based on BA classification. However, some control protocol packets will not be classified using BA classification, because they might not have a VLAN, MPLS, or IP header. These are:

- Untagged ARP packets
- Untagged Layer 2 control packets such as LACP or Ethernet OAM
- Untagged IS-IS packets

When the control queue feature is disabled, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to the restricted queue corresponding to the forwarding class associated with queue 0, as shown in the following two examples.

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the forwarding class (FC) associated with queue 0 is "be" (based on the **forwarding-class** statement configuration). "be" maps to restricted-queue number 3 (based on the "restricted-queue" configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to ingress queue 3 (not to ingress queue 0).

```
[edit chassis]
forwarding-classes {
  queue 0 be;
  queue 1 af-low8;
  queue 2 af-high;
  queue 3 ef;
  queue 4 ops_control;
  queue 5 net_control;
  queue 6 af-low10_12;
  }
  restricted-queues {
    forwarding-class ef queue-num 0;
    forwarding-class af-low8 queue-num 1;
    forwarding-class af-low10_12 queue-num 1;
    forwarding-class af-high queue-num 2;
    forwarding-class be queue-num 3;
```



```
}
```

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the FC associated with queue 0 is "ef" (based on the **forwarding-class** statement configuration). "ef" maps to restricted-queue number 0 (based on the **restricted-queue** statement configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets would go to ingress queue 0.

For tagged ARP, IS-IS, or Layer2 control packets, users should configure an explicit dot1p/dot1ad classifier to make sure these packets are directed to the correct queue. Without an explicit dot1p/dot1ad classifier, tagged ARP, IS-IS, or Layer 2 control packets will go to the restricted-queue corresponding to the forwarding class associated with queue 0.

```
[edit chassis]
forwarding-classes {
  queue 0 ef; <<< ef and be are interchanged
  queue 1 af-low8;
  queue 2 af-high;
  queue 3 be; <<< ef and be are interchanged
  queue 4 ops_control;
  queue 5 net_control;
  queue 6 af-low10_12;
  }
  restricted-queues {
    forwarding-class ef queue-num 0;
    forwarding-class af-low8 queue-num 1;
    forwarding-class af-low10_12 queue-num 1;
    forwarding-class af-high queue-num 2;
    forwarding-class be queue-num 3;
  }
```

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 395
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 404
- **no-pre-classifier on page 1278**
- Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC on page 407

- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC

Table 32 on page 408 lists the scenarios of handling oversubscription on the 10-port 10-Gigabit Ethernet LAN/WAN PIC for different combinations of port groups and active ports on the PIC.

Table 32: Handling Oversubscription on 10-Gigabit Ethernet LAN/WAN PICs

Number of Port Groups with Two Active Ports (A)	Number of Port Groups with One Active Port (B)	Total Number of Ports Used on PIC (C = Ax2 + B)	Status of Oversubscription and Throughput
0	1	1	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	2	2	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	5	5	Oversubscription is not active. Each port will receive 10 Gbps throughput.
1	0	2	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
1	4	6	Oversubscription is active for the port group that has two active ports. Each port in this port group will receive 5 Gbps throughput (with default shaper configuration). For the remaining four ports, oversubscription is not active. Each port will receive 10 Gbps throughput.
3	0	6	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
5	0	10	Oversubscription is active on all 10 ports (5 port groups). Each port will receive 5 Gbps throughput (with default shaper configuration).

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 395
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 404
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 404
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Mixed-Rate Mode Operation

To configure mixed-rate mode operation for a PF-24XGE-SFPP PIC:

1. Navigate to the **[edit chassis]** hierarchy level.
2. On a T4000 router, configure the mixed-rate mode by including the **mixed-rate-mode** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level.

[edit chassis]
user@host# **set fpc *fpc-slot* pic pic-number mixed-rate-mode**

On an LCC in a routing matrix, configure the mixed-rate mode by including the **mixed-rate-mode** statement at the **[edit chassis lcc *lcc number* fpc *slot-number* pic *pic-number*]** hierarchy level.

[edit chassis]
user@host# **set lcc *lcc number* fpc *fpc-slot* pic pic-number mixed-rate-mode**

3. Specify the port and the port speed that need to be configured. You can use one of the following speed attributes for this configuration.

   [edit chassis]
   user@host# **set fpc *fpc-slot* pic *pic-number* port port-number speed 1G;**
   user@host# **set fpc *fpc-slot* pic *pic-number* port port-number speed 10G;**
   user@host# **set lcc *lcc number* fpc *fpc-slot* pic pic-number speed 1G;**
   user@host# **set lcc *lcc number* fpc *fpc-slot* pic pic-number speed 10G;**

   NOTE: On a 12 port 10-Gigabit Ethernet PIC (PF-12XGE-SFPP), you can configure the port speed as 1G by including the **set fpc *fpc-slot* pic *pic-number* port port-number speed 1G** statement at the **[edit chassis]** hierarchy level.

   NOTE: To change the port speed from 10 Gbps to 1 Gbps on PF-24XGE-SFPP and PF-12XGE-SFPP PICs, SFP optics is required.

To disable mixed-rate mode operation, include the **delete chassis fpc *x* pic *y* mixed-rate-mode** statement at the **[edit chassis]** hierarchy level.

Related Documentation

- Modes of Operation of 10-Gigabit Ethernet PICs on page 403
- mixed-rate-mode on page 1261

P2-10G-40G-QSFPP PIC Overview

Starting with Junos OS Release 14.1R2 and 14.2R1, the PTX5000 Packet Transport Router supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC.

All the ports on the P2-10G-40G-QSFPP PIC are plugged into quad small form-factor pluggable plus transceivers (QSFP+) that, in turn, are connected to fiber-optic cables that support both 10-Gigabit Ethernet standards and 40-Gigabit Ethernet standards, thereby enabling you to configure the PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode.

Starting from Junos OS Release 14.2R3 and 16.1R1, you can configure the ports on the PIC in 10-Gigabit Ethernet mode or 40-Gigabit Ethernet mode at the port group level.

The following sections describe the P2-10G-40G-QSFPP PIC and the various framing modes that are supported on it:

Understanding Dual Configuration on P2-10G-40G-QSFPP PIC

All the ports on the P2-10G-40G-QSFPP PIC are QSFP+ based—that is, all the ports are connected to fiber-optic cables by means of QSFP+ transceivers.

The QSFP+ module—which includes the transceiver and the fiber-optic cable—supports the following standards on the P2-10G-40G-QSFPP PIC:

- 10-Gigabit Ethernet in LAN PHY framing mode (also known as native Ethernet mode) and WAN PHY framing mode.

  Note that the ports follow a 4-level interface-naming convention—et-*fpc/pic/QSFP+ port*:*channel* in this mode.

- 40-Gigabit Ethernet in LAN PHY framing mode.

  Note that the ports follow a 3-level interface-naming convention—et-*fpc/pic/QSFP+ port* in this mode.

NOTE: **The P2-10G-40G-QSFPP PIC provides forty-eight 10-Gigabit Ethernet ports or twelve 40-Gigabit Ethernet ports. or .**

The PIC can be configured either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode with the **set chassis fpc *fpc-number* pic pic-number pic-mode (10G | 40G)** configuration command. By default, the PIC is configured in 10-Gigabit Ethernet LAN PHY framing mode.

NOTE:
If you want configure the PIC in 10-Gigabit Ethernet mode to operate in 40-Gigabit Ethernet mode, you must:

1. Delete all the interfaces in the PIC at the **[edit interfaces]** hierarchy level.
2. Configure the PIC to operate in 40-Gigabit Ethernet mode by using the **set chassis fpc *fpc-slot* pic *pic-slot* pic-mode 40G** configuration command and commit.

The PIC reboots and starts operating in the new mode.

The same procedure is applicable when you can configure the PIC in 40-Gigabit Ethernet PIC to operate in 10-Gigabit Ethernet mode. In this case, you must execute the **set chassis fpc *fpc-slot* pic *pic-slot* pic-mode 10G** configuration mode command.

To check the current diagnostics of the PIC, you must run the relevant operational mode CLI commands such as **show chassis hardware**, **show interfaces diagnostics optics *interface-name***,

Understanding Port Group

The FPC2-PTX-P1A FPC on PTX5000 routers can host two PICs and has eight Packet Forwarding Engines. The first four Packet Forwarding Engines on the FPC are associated with PIC 0 and the next four are associated with PIC 1.

All ports associated to one Packet Forwarding Engine compose a port group. Each PIC supports four Packet Forwarding Engines. Therefore, four port groups exist for each P2-10G-40G-QSFPP PIC.

Each Packet Forwarding Engine providesthroughput of 120 Gbps.

Points to Remember

Consider the following points when configuring the PIC at the port group level:

- You can configure the ports in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode at the port group level.
- You can configure the port speed only on the first port in the port group. That is, you must configure the port speed for the port group on the ports numbered 0, 3, 6, and 9—the first ports in the respective port groups. An error message is logged when you try to configure the speed on any other port in the port group and this configuration will not have any effect on the PIC.
- You can configure the port speed of a port group only when the value of the **pic-mode** statement at the **[edit chassis fpc *fpc-slot* pic *pic-slot* pic-mode]** is set to 10G or when the statement is not configured.
- You cannot configure different speeds for the ports in the same port group.
- You can configure different speeds for different port groups.

Port Group in 10-Gigabit Ethernet Mode

Each Packet Forwarding Engine supports twelve 10-Gigabit Ethernet ports in LAN PHY or in WAN PHY framing mode.

Note that when a port group is configured from 10-Gigabit Ethernet mode to 40-Gigabit Ethernet mode, the ports with 4-level interface-naming convention are deleted and three 40-Gigabit Ethernet mode ports with 3-level interface-naming convention are created.

Note that when the configuration of a port group is changed from 10-Gigabit Ethernet mode to 40-Gigabit Ethernet mode, the configuration of the twelve 10-Gigabit Ethernet ports is deleted and the 4-level interface-naming convention of the ports is also lost. Instead, three 40-Gigabit Ethernet ports are configured and these ports adhere to the 3-level interface-naming convention

Port Group in 40-Gigabit Ethernet Mode

Each Packet Forwarding Engine supports three 40-Gigabit Ethernet ports in LAN PHY framing mode.

Note that when the configuration of a port group is changed from 40-Gigabit Ethernet mode to 10-Gigabit Ethernet mode, the configuration of the three 40-Gigabit Ethernet ports is deleted and the 3-level interface-naming convention of the ports is also lost. Instead, twelve 10-Gigabit Ethernet ports are configured and these ports adhere to the 4-level interface-naming convention.

Port Number Mapping When Port Groups Are Configured

Table 33 on page 413shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode at the port group level.

Table 33: Port Number Mapping When Port Groups Are Configured

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
0 (0)	et-1/1/0	et-1/1/0:0
		et-1/1/0:1
		et-1/1/0:2
		et-1/1/0:3
	et-1/1/1	et-1/1/1:0
		et-1/1/1:1
		et-1/1/1:2
		et-1/1/1:3
	et-1/1/2	et-1/1/2:0
		et-1/1/2:1
		et-1/1/2:2
		et-1/1/2:3
3(1)	et-1/1/3	et-1/1/3:0
		et-1/1/3:1
		et-1/1/3:2
		et-1/1/3:3
	et-1/1/4	et-1/1/4:0
		et-1/1/4:1
		et-1/1/4:2
		et-1/1/4:3
	et-1/1/5	et-1/1/5:0
		et-1/1/5:1
		et-1/1/5:2
		et-1/1/5:3

Table 33: Port Number Mapping When Port Groups Are Configured (continued)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
6(2)	et-1/1/6	et-1/1/6:0 et-1/1/6:1 et-1/1/6:2 et-1/1/6:3
	et-1/1/7	et-1/1/7:0 et-1/1/7:1 et-1/1/7:2 et-1/1/7:3
	et-1/1/8	et-1/1/8:0 et-1/1/8:1 et-1/1/8:2 et-1/1/8:3
6(2)	et-1/1/6	et-1/1/6:0 et-1/1/6:1 et-1/1/6:2 et-1/1/6:3
	et-1/1/7	et-1/1/7:0 et-1/1/7:1 et-1/1/7:2 et-1/1/7:3
	et-1/1/8	et-1/1/8:0 et-1/1/8:1 et-1/1/8:2 et-1/1/8:3

Table 33: Port Number Mapping When Port Groups Are Configured (continued)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
9(3)	et-1/1/9	et-1/1/9:0 et-1/1/9:1 et-1/1/9:2 et-1/1/9:3
	et-1/1/10	et-1/1/10:0 et-1/1/10:1 et-1/1/10:2 et-1/1/10:3
	et-1/1/11	et-1/1/11:0 et-1/1/11:1 et-1/1/11:2 et-1/1/11:3

Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured

Table 34 on page 415 shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode when port groups are *not* configured on the P2-10G-40G-QSFPP PIC.

Table 34: Port Number Mapping When Port Groups Are Not Configured

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
0	et-1/1/0	et-1/1/0:0 et-1/1/0:1 et-1/1/0:2 et-1/1/0:3
1	et-1/1/1	et-1/1/1:0 et-1/1/1:1 et-1/1/1:2 et-1/1/1:3

Table 34: Port Number Mapping When Port Groups Are Not Configured (continued)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
2	et-1/1/2	et-1/1/2:0 et-1/1/2:1 et-1/1/2:2 et-1/1/2:3
3	et-1/1/3	et-1/1/3:0 et-1/1/3:1 et-1/1/3:2 et-1/1/3:3
4	et-1/1/4	et-1/1/4:0 et-1/1/4:1 et-1/1/4:2 et-1/1/4:3
5	et-1/1/5	et-1/1/5:0 et-1/1/5:1 et-1/1/5:2 et-1/1/5:3
6	et-1/1/6	et-1/1/6:0 et-1/1/6:1 et-1/1/6:2 et-1/1/6:3
7	et-1/1/7	et-1/1/7:0 et-1/1/7:1 et-1/1/7:2 et-1/1/7:3
8	et-1/1/8	et-1/1/8:0 et-1/1/8:1 et-1/1/8:2 et-1/1/8:3

Table 34: Port Number Mapping When Port Groups Are Not Configured (continued)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
9	et-1/1/9	et-1/1/9:0 et-1/1/9:1 et-1/1/9:2 et-1/1/9:3
10	et-1/1/10	et-1/1/10:0 et-1/1/10:1 et-1/1/10:2 et-1/1/10:3
11	et-1/1/11	et-1/1/11:0 et-1/1/11:1 et-1/1/11:2 et-1/1/11:3

10-Gigabit Ethernet Mode

A 10-Gigabit Ethernet interface can operate in 10-Gigabit Ethernet LAN PHY framing mode or in 10-Gigabit Ethernet WAN PHY framing mode.

You can configure a 10-Gigabit Ethernet interface at the **[edit interface *interface-name* framing-mode (lan-phy | wan-phy)]** hierarchy level to operate in 10-Gigabit Ethernet LAN PHY framing mode or in 10-Gigabit Ethernet WAN PHY framing mode.

Each P2-10G-40G-QSFPP PIC provides 48 physical interfaces. The interfaces are represented by the 4-level interface-naming convention—et-*fpc/pic/QSFP+ port*:*channel*, where the value of the *QSFP+ port* option ranges from 0 through 11 and the value of the *channel* option ranges from 0 through 3.

- Framing Mode Overview on page 417
- Supported Features on LAN PHY and WAN PHY Framing Mode on page 418

Framing Mode Overview

When a P2-10G-40G-QSFPP PIC is configured in 10-Gigabit Ethernet framing mode, it can operate in one of the following framing modes:

- LAN PHY framing mode. Note that by default, the PIC is in 10-Gigabit Ethernet LAN PHY framing mode. You can configure loopback at the **[edit interfaces *interface-name* sonet-options loopback]** hierarchy level.

NOTE: The ports are set to LAN PHY framing mode by default when the **framing-mode** statement is not configured at the **[edit interface *interface-name*]** hierarchy level.

- WAN PHY framing mode

Supported Features on LAN PHY and WAN PHY Framing Mode

The following features are supported in LAN PHY and WAN PHY framing mode when the PIC operates in 10-Gigabit Ethernet mode:

- The following are supported for WAN interface sublayer statistics, defects, and alarms when the PIC operates in WAN PHY framing mode:
  - GR 253 standard.
  - **show interfaces** ***interfaces-name*** operational mode command displays WAN interface sublayer statistics, defects and alarms.
  - Interrupt-driven notification for WAN interface sublayer defects.
  - Path trace and trigger options for WAN interface sublayer alarms.
  - Transmitting and receiving J1 (path trace) messages—J1 is a part of path overhead in a WAN interface sublayer frame.
- Line loopback and local loopback. Loopback is configured at the **[edit interfaces *interface-name* sonet-options loopback]** hierarchy level in WAN PHY framing mode.
- The defects PHY LOL (loss of light) and PHY PLL (loss of PLL lock) are detected and reported at the physical level in WAN PHY framing mode.

Fast reroute (FRR) in WAN PHY framing mode:

- Enable or disable preemptive fast reroute (FRR) options at the **[edit interfaces *interface-name* otn-options preemptive-fast-reroute]** hierarchy level.
- Configure thresholds and interval for the optical channel data unit (ODU) signal degradation (**odu-signal-degrade**) and the configurable pre-FEC bit error rate (BER) (**ber-threshold-signal-degrade**) at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level and the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level, respectively.

40-Gigabit Ethernet Mode

You can configure twelve 40-Gigabit Ethernet interfaces that operate in LAN PHY framing mode. The interfaces are represented by the 3-level interface-naming convention et-*fpc/pic/QSFP+ port*, where the value of the *QSFP+ port* variable ranges from 0 through 11.

Release History Table

Release	Description

Related Documentation

- Configuring the P2-10G-40G-QSFPP PIC on page 419

Configuring the P2-10G-40G-QSFPP PIC

Starting with Junos OS Release 14.1R2, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC. You can configure the P2-10G-40G-QSFPP PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode..

The following tasks explain how to configure the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode and to configure the framing modes on it.

Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode:

1. In configuration mode, go to the **[edit chassis]** hierarchy level.

   ```
   [edit]
   user@host# edit chassis
   ```

2. Configure the PIC in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode after specifying the required FPC slot and PIC slot. Note that all the PIC ports in a PIC are configured at once with this configuration command.

   ```
   [edit chassis]
   user@host# set fpc fpc-slot pic pic-slot pic-mode (10G | 40G)
   ```

Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC that is configured in 10-Gigabit Ethernet mode to operate in 40-Gigabit Ethernet mode:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level.

   ```
   [edit]
   ```

user@host# **edit interfaces**

2. Delete all the interfaces in the PIC, commit, and then move to the top of the hierarchy level.

   [edit interfaces]
   user@host# **delete** ***interface-name***
   user@host# **delete ...**
   user@host# **commit**
   user@host# **top**

3. Configure the PIC to operate in 40-Gigabit Ethernet mode and commit.

   [edit]
   user@host# **set chassis fpc** ***fpc-slot*** **pic** ***pic-slot*** **pic-mode 40G**
   user@host# **commit**

After the configuration is committed, the PIC reboots and starts operating in the 40-Gigabit Ethernet mode. You can now configure the parameters, such as encapsulation, framing mode, and so on, for the twelve 40-Gigabit Ethernet interfaces in the PIC as needed.

Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC that is configured in 40-Gigabit Ethernet mode to operate in 10-Gigabit Ethernet mode:

1. In configuration mode, go to the **[edit interfaces** ***interfaces-name*****]** hierarchy level.

   [edit]
   user@host# **edit interfaces** ***interface-name***

2. Delete all the interfaces in the PIC, commit, and then move to the top of the hierarchy level.

   [edit interfaces]
   user@host# **delete** ***interface-name***
   user@host# **delete ...**
   user@host# **commit**
   user@host# **top**

3. Configure the PIC to operate in 10-Gigabit Ethernet mode and commit.

   [edit]
   user@host# **set chassis fpc** ***fpc-slot*** **pic** ***pic-slot*** **pic-mode 10G**
   user@host# **commit**

After the configuration is committed, the PIC reboots and starts operating in the 10-Gigabit Ethernet mode. You can now configure the parameters, such as encapsulation, framing mode, and so on, for the forty-eight 10-Gigabit Ethernet interfaces in the PIC as needed.

Configuring the PIC at Port Group Level

Before You Begin

Verify that the **pic-mode** statement at the **[edit chassis fpc *fpc-slot* pic *pic-slot* pic-mode]** is not configured or that its value is set to 10G.

To configure a port group in the P2-10G-40G-QSFPP PIC to operate in 10-Gigabit Ethernet mode or 40-Gigabit Ethernet mode:

1. In configuration mode, go to the **[edit chassis fpc *fpc-slot* pic *pic-slot*]** hierarchy level.

   [edit]
   user@host# **edit chassis fpc *fpc-slot* pic *pic-slot***

2. Configure the port number as 0, 3, 6, or 9 and the speed as 10G or 40G. Note that you can configure the port speed only on the first port in the port group. That is, configure the port speed only on the ports numbered 0, 3, 6, and 9. An error message is displayed when you try to configure the speed on any other port in the port group.

   [edit chassis fpc *fpc-slot* pic *pic-slot*]
   user@host# **set port *port-number* speed (10G | 40G)**

NOTE: A system log message is logged when you try to configure a different port speed on a port when the port group is operating at another speed.

Configuring Framing Mode on P2-10G-40G-QSFPP PIC

You can configure LAN PHY, or WAN PHY framing mode when the PIC is operating in 10-Gigabit Ethernet mode. You can configure LAN PHY framing mode when the PIC is operating in 40-Gigabit Etherent mode. The following tasks explain how to configure the various framing modes on the PIC:

- Configuring LAN PHY or WAN PHY Framing Mode in 10-Gigabit Ethernet Mode on page 421
- Configuring LAN PHY Framing Mode in 40-Gigabit Ethernet Mode on page 422

Configuring LAN PHY or WAN PHY Framing Mode in 10-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode to operate in LAN PHY framing mode or in WAN PHY framing mode, you must configure the framing mode individually on all the interfaces:

1. In configuration mode, go to the **[edit interfaces *interfaces-name*]** hierarchy level, where the interface name is in et-*fpc/pic/port*:*channel* format.

   [edit]
   user@host# **edit interfaces *interface-name***

2. Configure the framing mode as LAN PHY or WAN PHY and commit.

[edit interfaces *interface-name*]
user@host# **set framing (lan-phy | wan-phy)**
user@host# **commit**

For example, you can configure the framing mode as LAN PHY or WAN PHY on the et-1/1/1:0 interface.

Configuring LAN PHY Framing Mode in 40-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC in 40-Gigabit Ethernet mode to operate in LAN PHY framing mode:

1. In configuration mode, go to the **[edit interfaces** ***interfaces-name*****]** hierarchy level, where the interface name is in et-*fpc/pic/port* format.

   [edit]
   user@host# **edit interfaces** ***interface-name***

2. Configure the framing mode as LAN PHY and commit.

   [edit interfaces *interface-name*]
   user@host# **set framing (lan-phy)**
   user@host# **commit**

   For example, you can configure the framing mode as LAN PHY on the et-2/2/2 interface.

Release History Table

Release	Description
14.1R2	Starting with Junos OS Release 14.1R2, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC.

Related Documentation

• P2-10G-40G-QSFPP PIC Overview on page 409

Example: Configuring the P2-10G-40G-QSFPP PIC

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.1R2 or Junos OS Release 14.2 or later
- One PTX5000 router with P2-10G-40G-QSFPP PIC

Overview

Starting with Junos OS Release 14.1R2 and 14.2R1, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC.

All the ports on the P2-10G-40G-QSFPP PIC are QSFP+ based—that is, all the ports are connected to fiber-optic cables by means of QSFP+ transceivers. The P2-10G-40G-QSFPP PIC provides forty-eight 10-Gigabit Ethernet ports or twelve 40-Gigabit Ethernet ports.

The QSFP+ module—which includes the transceiver and the fiber-optic cable—supports the following standards on the P2-10G-40G-QSFPP PIC:

- 10-Gigabit Ethernet in LAN PHY framing mode (also known as native Ethernet mode) and WAN PHY framing mode.
- 40-Gigabit Ethernet in LAN PHY framing mode.

Configuration

To configure the P2-10G-40G-QSFPP PIC to operate in 10-Gigabit Ethernet mode, and to set the framing mode and other options on an interface on this PIC, perform the following tasks:

- Configuring the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet Mode on page 423
- Configuring the Framing Mode on an Interface on page 423
- Configuring the Interface Options on page 424
- Verification on page 424

Configuring the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet Mode

Step-by-Step Procedure

Configure the PIC in 10-Gigabit Ethernet mode.

1. In configuration mode, go to the **[edit chassis]** hierarchy level.

   ```
   [edit]
   user@host# edit chassis
   ```

2. Configure the PIC in 10-Gigabit Ethernet mode after specifying the required FPC slot and PIC slot. Note that the PIC restarts after the configuration is committed and all the ports in the PIC come up in the 10-Gigabit Ethernet mode.

   ```
   [edit chassis]
   user@host# set fpc 1 pic 1 pic-mode 10G
   ```

Configuring the Framing Mode on an Interface

Step-by-Step Procedure

To configure an interface *et-1/1/1:0* in the P2-10G-40G-QSFPP PIC to operate in LAN PHY framing mode:

1. In configuration mode, go to the **[edit interfaces et-1/1/1:0]** hierarchy level.

```
[edit]
user@host# edit interfaces et-1/1/1:0
```

2. Configure the framing mode for the interface as LAN PHY and commit.

   ```
   [edit interfaces et-1/1/1:0]
   user@host# set framing lan-phy
   user@host# commit
   ```

   Similarly, you can configure LAN PHY or WAN PHY framing mode for the other interfaces in the PIC.

Configuring the Interface Options

Step-by-Step Procedure

Configure the interface options for the interface et-1/1/1:0 as needed. The following procedure configures a few interface-specific options.

1. In configuration mode, go to the **[edit interfaces et-1/1/1:0]** hierarchy level.

   ```
   [edit]
   user@host# edit interfaces et-1/1/1:0
   ```

2. Configure the encapsulation as ethernet-ccc.

   ```
   [edit interfaces et-1/1/1:0]
   user@host# set encapsulation ethernet-ccc
   ```

3. Configure the family as CCC for the logical interface 0.

   ```
   [edit interfaces et-1/1/1:0]
   user@host# set unit 0 family ccc
   ```

4. Enable flow control to regulate the flow of packets from the router to the remote side of the network connection.

   ```
   [edit interfaces et-1/1/1:0 gigether-options]
   user@host# set flow-control
   ```

5. Enable loopback mode for the interface, commit the configuration, and exit the configuration mode.

   ```
   [edit interfaces et-1/1/1:0 gigether-options]
   user@host# set loopback
   user@host# commit
   user@host# quit
   ```

Verification

Displaying Interface Details

Purpose To display interface-specific details of the et-1/1/1:0 interface.

Action Execute the **show interfaces et-1/1/1:0** operational command.

```
user@host# run show interfaces et-1/1/1:0
Interface index: 525, SNMP ifIndex: 522
  Link-level type: Ethernet, MTU: 1514, MRU: 0, LAN-PHY mode, Speed: 10Gbps, BPDU
 Error:
None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: ac:4b:c8:f6:af:68, Hardware address: ac:4b:c8:f6:af:68
  Last flapped   : 2014-07-25 02:23:56 PDT (02:16:07 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         1
  Interface transmit statistics: Disabled
```

Meaning The interface details are displayed. Note that to display information for an interface in 10-Gigabit Ethernet mode for the P2-10G-40G-QSFPP PIC, you must use the et-*fpc*/*pic*/*port*:*channel* format.

Release History Table

Release	Description

Related Documentation

- P2-10G-40G-QSFPP PIC Overview on page 409

CHAPTER 23

Configuring 10-Gigabit Ethernet Framing

10-Gigabit Ethernet Framing Overview

The 10-Gigabit Ethernet interfaces support operation in two modes:

- 10GBASE-R, LAN Physical Layer Device (LAN PHY)
- 10GBASE-W, WAN Physical Layer Device (WAN PHY)

When the external interface is running in LAN PHY mode, it bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface. When the external interface is running in WAN PHY mode, it uses the WIS sublayer to transport 10-Gigabit Ethernet frames in an OC192c SONET payload.

WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600, T4000 and PTX Series Packet Transport routers only.

NOTE: **The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).**

Although the external interface provides a lower throughput when running in WAN PHY mode because of the extra SONET overhead, it can interoperate with SONET section or line level repeaters. This creates an advantage when the interface is used for long-distance, point-to-point 10-Gigabit Ethernet links. When the external interface is running in WAN PHY mode, some SONET options are supported. For information about SONET options supported on this interface, see *Configuring SONET Options for 10-Gigabit Ethernet Interfaces*.

NOTE: SONET or SDH framing mode configuration **framing (sdh | sonet)** is not applicable on the 10-Gigabit Ethernet ports. Configuring the **wan-phy** framing mode on the 10-Gigabit Ethernet ports allows the interface to accept SONET or SDH frames without further configuration.

Release History Table

Release	Description
12.2	Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
12.1R2	Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP).

Related Documentation

- *Configuring SONET/SDH Framing Mode for Ports*
- Configuring 10-Gigabit Ethernet Framing on page 429
- Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces on page 428
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces

If you use the **wan-phy** statement option at the **[edit interfaces xe-*fpc*/*pic*/0 framing]** hierarchy level to configure Trio WAN mode framing for 10-Gigabit Ethernet interfaces, then the alarm behavior of the link, although in full compliance with the IEEE 802.3ae 10-Gigabit Ethernet standard, might not be as expected.

In particular:

- The interface does not distinguish between loss of light (LOL), loss of phase lock loop (PLL), or loss of signal (LOS). If a loss of PLL or LOS alarm occurs, then both PLL and LOS alarms are raised. LOL is also raised because there is no separate LOL indication from the hardware.
- The interface does not raise LOS, PLL, or LOL alarms when the fiber in disconnected from the interface port. You must remove the hardware to raise this alarm.
- The interface line-level alarm indicator signal (AIS-L) is not always raised in response to a loss of framing (LOF) defect alarm.
- If the AIS-L or path-level AIS (AIS-P) occurs, the interface path-level loss of code delineation (LCD-P) is not detected. LCD-P is seen during the path-level remote defect indicator (RDI-P) alarm.
- If an AIS-L alarm occurs, the AIS-P is not detected, but the LOP alarm is detected.

None of the alarm issues are misleading, but they make troubleshooting the root cause of problems more complex.

Related Documentation

- **framing on page 1179**
- Configuring 10-Gigabit Ethernet Framing on page 429
- 10-Gigabit Ethernet Framing Overview on page 427

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring 10-Gigabit Ethernet Framing

The 10-Gigabit Ethernet interfaces uses the interface type **xe-*fpc/pic/port***. On single port devices, the port number is always zero.

The **xe-*fpc/pic/port*** interface inherits all the configuration commands that are used for gigabit Ethernet (**ge-*fpc/pic/port***) interfaces.

To configure LAN PHY or WAN PHY operating mode, include the **framing** statement with the **lan-phy** or **wan-phy** option at the **[edit interfaces xe-*fpc* /*pic*/0]** hierarchy level.

```
[edit interfaces xe-fpc/pic/0 framing]
framing (lan-phy | wan-phy);
```

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- On PTX Series Transport Routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.

To display interface information, use the operational mode command **show interfaces xe-*fpc/pic/port* extensive**.

NOTE:

- SONET or SDH framing mode configuration **framing (sdh | sonet)** is not applicable on the 10-Gigabit Ethernet ports. Configuring the **wan-phy** framing mode on the 10-Gigabit Ethernet ports allows the interface to accept SONET or SDH frames without further configuration.
- If you configure the WAN PHY mode on an aggregated Ethernet interface, you must set the aggregated Ethernet link speed to OC192.

Related Documentation

- framing on page 1179
- 10-Gigabit Ethernet Framing Overview on page 427
- Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces on page 428

- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 24

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

Gigabit Ethernet Notification of Link Down Alarm Overview

Notification of link down alarm generation and transfer is supported for all 10-Gigabit Ethernet PIC interfaces on M120 and M320 routers. On the MX Series and T series routers, notification of link down alarm generation and transfer is supported for all Gigabit Ethernet Interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit).

Related Documentation

- Configuring Gigabit Ethernet Notification of Link Down Alarm on page 432
- **asynchronous-notification on page 1084**

- *Ethernet Interfaces Feature Guide for Routing Devices*

10-Gigabit Ethernet Notification of Link Down for Optics Options Overview

Notification of link down is supported for IQ2 10-Gigabit Ethernet interfaces and MX Series DPCs. You can use link down notification to help identify optical link connectivity problems.

For information on configuring link down notification, see “Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning” on page 432.

Related Documentation

- Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning on page 432

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gigabit Ethernet Notification of Link Down Alarm

Notification of link down alarm generation and transfer is supported for all 10-Gigabit Ethernet PIC interfaces on M120 and M320 routers. On the MX Series and T Series routers, notification of link down alarm generation and transfer is supported for all Gigabit Ethernet Interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit).

To configure this option, include the **asynchronous-notification** statement at the **[edit interfaces ge-** ***fpc/pic/port*** **gigether-options]** hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
  gigether-options {
    asynchronous-notification;
  }
}
```

Related Documentation

- Gigabit Ethernet Notification of Link Down Alarm Overview on page 431
- *Ethernet Interfaces Feature Guide for Routing Devices*
- asynchronous-notification on page 1084

Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning

To configure this option, include the **alarm** or **warning** statement at the **[edit interfaces ge-** ***fpc/pic/port*** **optics-options]** hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
  optics-options {
    alarm alarm-name {
      (syslog | link-down);
      }
    warning warning-name {
      (syslog | link-down);
      }
  }
}
```

Related Documentation

- alarm on page 938
- warning on page 984
- 10-Gigabit Ethernet Notification of Link Down for Optics Options Overview on page 431
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 25

Configuring 40-Gigabit Ethernet PICs

40-Gigabit Ethernet PIC Overview

The 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP) is a 1-port 40-Gigabit Ethernet Type 4 PIC with C form-factor pluggable transceiver (CFP) optics supported on T640, T1600, and T4000 routers. The 40-Gigabit Ethernet PIC occupies FPC slot 0 or 1 in the Type 4 FPC and it is similar to any regular PIC such as the 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP) PIC. The CFP information appears under the PIC information in the show command output.

The 40-Gigabit Ethernet PIC with CFP supports flexible Ethernet services encapsulation and MAC accounting.

MAC learning, MAC policing, and Layer 2 rewrite features are not supported.

The 40-Gigabit Ethernet PIC with CFP supports the following features:

- Encapsulation protocols such as:
  - Layer 2 protocols
    - Ethernet CCC, Ethernet TCC, and Ethernet VPLS
    - VLAN CCC
    - Extended VLAN TCC
    - VLAN VPLS
    - Flexible Ethernet service
  - Layer 3 protocols
    - IPv4
    - IPv6

    - MPLS

- CFP Multi-Source Agreement (MSA)-compliant management data input/output (MDIO) control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) (in all PIC and chassis configurations).
- Interface creation:
    - When the PIC is brought online, the router creates one interface, et-*x*/*y*/0, where *x* represents the FPC slot number and *y* represents PIC slot number. The physical interface represents internal Ethernet Packet Forwarding Engines.
    - The FPC slot number ranges from 0 through 7 in T640, T1600, and T4000 routers. The PIC slot numbers are 0 and 1.
    - Packet Forwarding Engine 0 is the physical interface 0, and Packet Forwarding Engine 1 is the physical interface 1.
- 802.3 link aggregation:
    - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
    - An aggregate bundle that consists purely of 40-Gigabit Ethernet PICs supports a maximum of 40-Gigabit Ethernet links depending on the system implementation.

For Junos OS configuration information about this PIC, see “Configuring 40-Gigabit Ethernet PICs” on page 435. For hardware compatibility information, see the *T1600 PICs Supported* topic in the *T1600 Core Router Hardware Guide* hardware guide and the *T640 PICs Supported* topic in the *T640 Core Router Hardware Guide* hardware guide, and the *T4000 PICs Supported* topic in the *T4000 Core Router Hardware Guide* hardware guide.

Related Documentation

- Configuring 40-Gigabit Ethernet PICs on page 435
- *T640 Core Router Hardware Guide*
- *T1600 Core Router Hardware Guide*
- *T4000 Core Router Hardware Guide*
- *TX Matrix Plus Router Hardware Guide*
- *T640 PICs Supported*
- *T1600 PICs Supported*
- *T4000 PICs Supported*

Configuring 40-Gigabit Ethernet PICs

You can configure the following features on the 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP):

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting for receive (**Rx**) and transmit (**Tx**)
- Multiple tag protocol ID (TPID) support
- Channels defined by two stacked VLAN tags
- Channels defined by **flex-vlan-tagging**
- IP service for stacked VLAN tags
- IP service for nonstandard TPID

The following features are not supported on the 40-Gigabit Ethernet PIC with CFP:

- MAC learning
- MAC policing
- Layer 2 rewrite

NOTE: Each 40-Gigabit Ethernet PIC with CFP creates a single et- physical interface in the Routing Engine and Packet Forwarding Engine.

The 40-Gigabit Ethernet PIC with CFP supports aggregated Ethernet configuration to achieve higher throughput capability, whereby the configuration is similar to the 1-Gigabit or 10-Gigabit aggregated Ethernet interface configuration. A maximum of 40-Gigabit Ethernet PIC links can be bundled into a single aggregated Ethernet configuration depending on the system implementation.

To configure the 40-Gigabit Ethernet PIC with CFP:

1. Perform the media configuration.

   The command used to configure the media for the 40-Gigabit Ethernet PIC with CFP is the same as that for other Ethernet PICs, such as the 4-port 10-Gigabit Ethernet PIC.

2. Specify the logical interfaces.

A single physical interface is created when the 40-Gigabit Ethernet PIC with CFP is brought online (et-*x*/*y*/0, where *x* represents the FPC slot number and *y* represents the PIC slot number). For more information, see “Configuring Access Mode on a Logical Interface” on page 261 and “Configuring a Logical Interface for Trunk Mode” on page 262.

3. Configure the 802.3 link aggregation.

   - You must explicitly configure an aggregated interface on the 40-Gigabit Ethernet PIC with CFP that includes the 40-Gigabit Ethernet interfaces. For more information, see “Configuring an Aggregated Ethernet Interface” on page 108.
   - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with the configuration of the 1-Gigabit Ethernet PIC, 10-Gigabit Ethernet PIC, and the aggregated Ethernet interfaces. In each aggregated bundle, Junos OS supports a maximum of 40-Gigabit Ethernet links. For more information, see “Configuring an Aggregated Ethernet Interface” on page 108 and “10-port 10-Gigabit Ethernet LAN/WAN PIC Overview” on page 395.

4. Configure the Packet Forwarding Engine features.

   The 40-Gigabit Ethernet PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see “Configuring Gigabit Ethernet Policers” on page 539, “Configuring MAC Address Filtering” on page 544, and “Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview” on page 559.

Related Documentation

- 40-Gigabit Ethernet PIC Overview on page 433
- Configuring Gigabit Ethernet Policers on page 539
- Configuring MAC Address Filtering on page 544
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559

CHAPTER 26

Configuring 100-Gigabit Ethernet PICs/MICs

100-Gigabit Ethernet Interfaces Overview

MX Series 100-Gigabit Ethernet Interfaces

Table 35 on page 438 lists the 100-Gigabit Ethernet interfaces supported by MX Series routers.

Table 35: MX Series 100-Gigabit Ethernet Interfaces

Interface Module	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet MIC with CFP	MIC3-3D-1X100GE-CFP	MX240 MX480 MX960 MX2010 MX2020	*100-Gigabit Ethernet MIC with CFP* "MPC3E MIC Overview" on page 440
100-Gigabit Ethernet MIC with CXP	MIC3-3D-1X100GE-CXP	MX240 MX480 MX960 MX2010 MX2020	*100-Gigabit Ethernet MIC with CXP* "MPC3E MIC Overview" on page 440
100-Gigabit Ethernet ports on the MPC4E	MPC4E-3D-2CGE-8XGE	MX240 MX480 MX960 MX2010 MX2020	*MPC4E on MX Series Routers Overview* *2x100GE + 8x10GE MPC4E*
100-Gigabit Ethernet MIC with CFP2	MIC6-100G-CFP2	MX2010 MX2020	*100-Gigabit Ethernet MIC with CFP2*
100-Gigabit Ethernet MIC with CXP (4 Ports)	MIC6-100G-CXP	MX2010 MX2020	*100-Gigabit Ethernet MIC with CXP (4 Ports)*

PTX Series 100-Gigabit Ethernet Interfaces

Table 36 on page 438 lists the 100-Gigabit Ethernet interfaces supported by PTX Series routers.

Table 36: PTX Series 100-Gigabit Ethernet Interfaces

PIC	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet PIC with CFP	P1-PTX-2-100GE-CFP	PTX5000	*100-Gigabit Ethernet PIC with CFP (PTX Series)*
100-Gigabit Ethernet PIC with CFP2	P2-100GE-CFP2	PTX5000	*100-Gigabit Ethernet PIC with CFP2 (PTX Series)*

Table 36: PTX Series 100-Gigabit Ethernet Interfaces (continued)

PIC	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet OTN PIC	P2-100GE-OTN	PTX5000	*100-Gigabit Ethernet OTN PIC with CFP2 (PTX Series)* "Understanding the P2-100GE-OTN PIC" on page 500 "Configuring OTN Interfaces on P2-100GE-OTN PIC" on page 504
100-Gigabit DWDM OTN PIC	P1-PTX-2-100G-WDM	PTX5000 PTX3000	*100-Gigabit DWDM OTN PIC (PTX Series)*

T Series 100-Gigabit Ethernet Interfaces

Table 37 on page 439 lists the 100-Gigabit Ethernet interfaces supported by T Series routers.

Table 37: T Series 100-Gigabit Ethernet Interfaces

PIC	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet PIC with CFP (Type 4)	PD-1CE-CFP-FPC4	T1600 T4000	*100-Gigabit Ethernet PIC with CFP (T1600 Router)* *100-Gigabit Ethernet PIC with CFP (T4000 Router)* "100-Gigabit Ethernet Type 4 PIC with CFP Overview" on page 441 "Configuring 100-Gigabit Ethernet Type 4 PIC With CFP" on page 444
100-Gigabit Ethernet PIC with CFP (Type 5)	PF-1CGE-CFP	T4000	*100-Gigabit Ethernet PIC with CFP (T4000 Router)* "100-Gigabit Ethernet Type 5 PIC with CFP Overview" on page 450

Related Documentation

- *MICs Supported by MX Series Routers*
- *MPCs Supported by MX Series Routers*
- *PICs Supported on the PTX Series*
- *T1600 PICs Supported*
- *T4000 PICs Supported*

MPC3E MIC Overview

The MPC3E supports two separate slots for MICs. MICs provide the physical interface and are installed into the MPCs.

The MPC3E supports these MICs as field replaceable units (FRUs):

- *100-Gigabit Ethernet MIC with CFP* (model number MIC3-3D-1X100GE-CFP)
- *100-Gigabit Ethernet MIC with CXP* (model number MIC3-3D-1X100GE-CXP)
- 10-port 10-Gigabit Ethernet MIC with SFPP (model number MIC3-3D-10XGE-SFPP)
- 2-port *40-Gigabit Ethernet MIC with QSFP+* (model number MIC3-3D-2X40GE-QSFPP)

The MPC3E has two separate configurable MIC slots. Each MIC corresponds to a single PIC and the mapping between the MIC and PIC is 1 to 1 (one MIC is treated as one PIC). The MIC plugged into slot 0 corresponds to PIC 0 and the MIC plugged into slot 1 corresponds to PIC 2.

The MPC3E also supports these legacy MICs:

- 20-port *Gigabit Ethernet MIC with SFP* (model number MIC-3D-20GE-SFP)
- 2-port *10-Gigabit Ethernet MICs with XFP* (model number MIC-3D-2XGE-XFP)

The 100-Gigabit Ethernet CFP MIC supports the IEEE standards—compliant 100BASE-LR4 interface, using the 100G CFP optical transceiver modules for connectivity. The 100-Gigabit Ethernet CXP MIC supports the 100BASE-SR10 interface, using 100-Gigabit CXP optical transceiver modules for connectivity. The 2-port 40-Gigabit Ethernet QSFPP MIC supports the 40BASE-SR4 interface and uses quad small form-factor pluggable (QSFPP) optical transceivers for connectivity. The 10-port 10-Gigabit Ethernet SFPP MIC uses SFP+ optical transceiver modules for connectivity.

For detailed information about each MIC, see *100-Gigabit Ethernet MIC with CFP*, *100-Gigabit Ethernet MIC with CXP*, *40-Gigabit Ethernet MIC with QSFP+*. For information about supported hardware and transceivers, see *MPC3E*.

The MPC3E supports these features:

- Optical diagnostics and related alarms
- Virtual Router Redundancy Protocol (VRRP) support
- IEEE 802.1Q virtual LANs (VLANs) support
- Synchronous Ethernet
- Remote monitoring (RMON) and Ethernet statistics (EtherStats)
- Source MAC learning
- MAC accounting and policing—Dynamic local address learning of source MAC addresses
- Flexible Ethernet encapsulation
- Multiple Tag Protocol Identifiers (TPIDs)

NOTE: The MPC3E supports Ethernet interfaces only. SONET interfaces are not supported.

For information about the supported and unsupported Junos OS features for this MPC, see “Protocols and Applications Supported by the MPC3E (MX-MPC3E)” in the *MX Series Interface Module Reference*.

Related Documentation

- *MPC3E on MX Series Routers Overview*
- *Protocols and Applications Supported by the MPC3E on MX Series Routers*
- *100-Gigabit Ethernet MIC with CFP*
- *100-Gigabit Ethernet MIC with CXP*
- 2-port*40-Gigabit Ethernet MIC with QSFP+*
- 2-port *10-Gigabit Ethernet MICs with XFP*
- *MX Series Interface Module Reference*

100-Gigabit Ethernet Type 4 PIC with CFP Overview

The 100-Gigabit Ethernet PIC (model number PD-1CE-CFP-FPC4) is a 1-port 100-Gigabit Ethernet Type 4 PIC with 100-gigabit small form-factor pluggable (CFP) transceiver. This PIC is available only as packaged in an assembly with the T1600-FPC4-ES FPC. The 100-Gigabit Ethernet PIC occupies PIC slots 0 and 1 in the T1600-FPC4-ES FPC. For information about supported transceivers and hardware, see *100-Gigabit Ethernet PIC with CFP (T1600 Router)*.

The 100-Gigabit Ethernet PIC supports flexible encapsulation and MAC accounting.

MAC learning, MAC policing, and Layer 2 rewrite functionality are not supported.

The ingress flow can be filtered based on the VLAN source and destination addresses. Ingress frames can also be classified according to VLAN, stacked VLAN, source address, VLAN source address, and stacked VLAN source address. VLAN manipulation on egress frames are supported on both outer and inner VLAN tags.

The following features are supported:

- The following encapsulation protocols are supported:
  - Layer 2 protocols
    - Ethernet CCC, Ethernet TCC, Ethernet VPLS
    - VLAN CCC
    - Extended VLAN TCC
    - VLAN VPLS
    - Flexible Ethernet service

- Layer 3 protocols
  - IPv4
  - Ipv6
  - MPLS
- CFP MSA compliant MDIO control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) is supported in all PIC and chassis configurations.
- Interface creation:
  - When the PIC, is brought online, the router creates two 50 gigabit capable interfaces, **et-*x*/0/0:0** and **et-*x*/0/0:1**, where *x* represents the FPC slot number. Each physical interface represents two internal 50 gigabit Ethernet Packet Forwarding Engines. Two logical interfaces are configured under each physical interface.
  - Packet Forwarding Engine 0 is physical interface 0, Packet Forwarding Engine 1 is physical interface 1
- 802.3 link aggregation:

  Same rate or same mode link aggregation:

  - Two logical interfaces are created for each 100-Gigabit Ethernet PIC. To utilize bandwidth beyond 50 gigabits per second, an aggregate interface must be explicitly configured on the 100-Gigabit Ethernet PIC that includes the two 50 gigabit interfaces.
  - Each 100 gigabit Ethernet aggregate consumes one of the router-wide aggregated Ethernet device pools. The number of 100-Gigabit Ethernet PICs cannot exceed the router-wide limit, which is 128 for Ethernet.
  - In each aggregate bundle, each 100-Gigabit Ethernet PIC consumes two members. Hence, an aggregate bundle that consists purely of 100-Gigabit Ethernet PICs supports a maximum of half of the software limit for the number of members. Therefore, with a maximum of 16 links, up to 8 100-Gigabit Ethernet links are supported.
  - Combining 100-Gigabit Ethernet PICs into aggregate interfaces with other Ethernet PICs is not permitted. However, other Ethernet PICs can also be configured within the same T1600 with 100-Gigabit Ethernet PICs, and used in separate aggregate interfaces.
  - Multiple (Juniper Networks) Type 4 100-Gigabit Ethernet PICs on a T1600 router can be combined into a static aggregated Ethernet bundle to connect to a different type of 100 gigabit Ethernet PIC on a remote router (Juniper Networks or other vendors). LACP is not supported in this configuration.

Mixed rate or mixed mode link aggregation:

- Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC.
- Static link protection and Link Aggregation Control Protocol (LACP) is supported on mixed aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC.
- When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC, ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC.

- Software Packet Forwarding Engine—Supports all Gigabit Ethernet PIC classification, firewall filter, queuing model, and rewrite functionality.
- Egress traffic performance—Maximum egress throughput is 100 gigabits per second on the physical interface, with 50 gigabits per second on the two assigned logical interfaces.
- Ingress traffic performance—Maximum ingress throughput is 100 gigabits per second on the physical interface, with 50 gigabits per second on the two assigned logical interfaces. To achieve 100 gigabits per second ingress traffic performance, use one of the interoperability modes described below. For example, if VLAN steering mode is not used when connecting to a remote 100 gigabits per second interface (that is on a different 100 gigabits per second PIC on a Juniper Networks router or a different vendor's equipment), then all ingress traffic will try to use one of the 50 gigabits per second Packet Forwarding Engines, rather than be distributed among the two 50 gigabits per second Packet Forwarding Engines, resulting in a total of 50 gigabits per second ingress performance.
- Interoperability modes—The 100-Gigabit Ethernet PIC supports interoperability with through configuration in one of the following two forwarding option modes:
  - *SA multicast mode*—In this mode, the 100-Gigabit Ethernet PIC supports interconnection with other Juniper Networks 100-Gigabit Ethernet PICs (Model: PD-1CE-CFP) interfaces only.
  - *VLAN steering mode*—In this mode, the 100-Gigabit Ethernet Type 4 PIC with CFP supports interoperability with 100 gigabit Ethernet interfaces from other vendors only.

Related Documentation

- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP on page 444
- *T1600 Core Router Hardware Guide*
- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

Configuring 100-Gigabit Ethernet Type 4 PIC With CFP

You can configure the following features on the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4):

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting in RX
- Channels defined by two stacked VLAN tags
- Channels defined by flex-vlan-tagging
- IP service for stacked VLAN tags
- Layer 2 rewrite

The following features are not supported on the 100-Gigabit Ethernet Type 4 PIC with CFP:

- Multiple TPID
- IP service for non-standard TPID
- MAC learning
- MAC policing

NOTE:

- For the 100-Gigabit Ethernet Type 4 PIC with CFP, only the PIC0 **online** and **offline** CLI commands are supported. The PIC1 **online** and **offline** CLI commands are not supported.
- Each 100-Gigabit Ethernet Type 4 PIC with CFP creates two et- physical interfaces, defined as 50-gigabit physical interfaces in the Routing Engine and Packet Forwarding Engine. By default, these are independent physical interfaces and are not configured as an aggregated Ethernet interface.

To configure a 100-Gigabit Ethernet Type 4 PIC with CFP:

1. Perform the media configuration:

   The 100-Gigabit Ethernet Type 4 PIC with CFP features a 100 gigabit per second pipe. The media-related configuration commands for **et-*x*/0/0:0** and **et-*x*/0/0:1** must both be configured at the same time and configured with the same value, otherwise the commit operation fails.

   When configuring to activate or deactivate the interface, if the interface contains the described media-related configuration, it must activate and deactivate both units 0 and 1 at the same time, otherwise the commit operation fails.

   The following media configuration commands have the above described restriction:

   - # **set interfaces et-x/0/0:1 disable**
   - # **set interfaces et-x/0/0:1 gigether-options loopback**
   - # **set interfaces et-x/0/0:1 mtu yyy**

   Due to an MTU restriction, the vlan-tagging and flexible-vlan-tagging configuration on et-x/0/0:0 and et-x/0/0:1 must be same, otherwise the commit operation fails.

2. Specify the logical interfaces:

   a. Two physical interfaces are created when the 100-Gigabit Ethernet Type 4 PIC with CFP is brought online (**et-*x*/0/0:0** and **et-*x*/0/0:1**, where *x* represents the FPC slot number). Each physical interface represents two internal 50-gigabit Ethernet Packet Forwarding Engines.

   b. Two logical interfaces are configured under each physical interface: Packet Forwarding Engine 0 is physical interface 0 and Packet Forwarding Engine 1 is physical interface 1.

3. Configure the 802.3 link aggregation:

   a. The 100-Gigabit Ethernet PIC supports aggregated Ethernet configuration to achieve higher throughput capability, whereby configuration is similar to the 1G/10G aggregated Ethernet interface configuration.

   b. Two physical interfaces are created for each 100-Gigabit Ethernet Type 4 PIC with CFP. To utilize bandwidth beyond 50 gigabits, a same rate and same mode aggregated Ethernet interface must be explicitly configured on the 100-Gigabit Ethernet Type 4 PIC with CFP that includes these two 50-gigabit interfaces.

   c. Each 100-Gigabit Ethernet Type 4 PIC with CFP aggregate consumes one of the router-wide aggregated Ethernet device pools. In Junos OS with 100-Gigabit Ethernet PICs, you cannot exceed the router limit of 128 Ethernet PICs.

   d. In each aggregated bundle, each 100-Gigabit Ethernet Type 4 PIC with CFP consumes two aggregate members. Hence, an aggregated bundle consisting of only one 100-Gigabit Ethernet Type 4 PIC with CFP supports only up to half of the Junos OS limit for the number of members. The Junos OS supports a maximum of 16 links for up to 8 100-Gigabit Ethernet Type 4 PIC with CFP links.

NOTE:
The 100-Gigabit Ethernet Type 4 PIC with CFP has the following restrictions for same rate and same mode aggregated Ethernet configuration:

- Both physical interfaces belonging to the same 100-Gigabit Ethernet PIC must be included in the same aggregated Ethernet physical interfaces. The aggregation of the 100-Gigabit Ethernet PIC interface is always an even number of physical interfaces.
- The 100-Gigabit Ethernet PIC physical interface cannot be configured in the aggregated interface with any other type of physical interface.
- The maximum supported number of aggregated 100-Gigabit Ethernet PIC interfaces is half of the number that the Junos OS supports for 1G/10G aggregated Ethernet. For example, if Junos OS supports 16 ports of 10-gigabit Ethernet aggregation, it supports 8 ports of 100-Gigabit Ethernet PIC aggregation. This is because each port of the 100-Gigabit Ethernet PIC port using 2 physical interfaces (et-x/0/0:0 and et-x/0/0:1), where each physical interface represents 50 gigabits of traffic capacity.

e. Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC. When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC, ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.

NOTE:

The 100-Gigabit Ethernet Type 4 PIC with CFP has the following restrictions for mixed rate and mixed mode aggregated Ethernet configuration:

- A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
- Traffic distribution is based on the hash calculated on the egress packet header. Hash range is fairly distributed according to member links' speed. This guarantees hash fairness but it does not guarantee fair traffic distribution depending on the rate of the egress streams.
- Packets are dropped when the total throughput of the hash flow exiting a member link (or multiple hash flows exiting a single member link) exceeds the link speed of the member link. This can happen when egress member link changes because of a link failure and the hash flow switches to a member link of speed that is less than the total throughput of the hash flow.
- Rate-based CoS components such as scheduler, shaper, and policer are not supported on mixed rate aggregated Ethernet links. However, the default CoS settings are supported by default on the mixed rate aggregated Ethernet links.
- Load balancing is performed at the ingress Packet Forwarding Engine. Therefore, you must ensure that the egress traffic on the aggregated Ethernet link enters through the hardware platforms that support mixed aggregated Ethernet bundles.
- Mixed aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed aggregated Ethernet load balancing is configured at egress.
- Load balancing of the egress traffic across the member links of a mixed rate aggregated Ethernet link is proportional to the rates of the member links.
- Egress multicast load balancing is not supported on mixed aggregated Ethernet interfaces.
- Changing the **edit interfaces ae*x* aggregated-ether-options link-speed** configuration of a mixed aggregated Ethernet link, which is configured on the supported interfaces of on T640, T1600, T4000, and TX Matrix Plus routers, leads to aggregated Ethernet link flapping.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC, changing aggregated Ethernet link protection configurations leads to aggregated Ethernet link flapping.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP.

- The **show interfaces ae*x*** command displays the link speed of the aggregated Ethernet interface, which is the sum of the link speeds of all the active member links.

4. Configure the Packet Forwarding Engine features:

   a. The 100-Gigabit Ethernet Type 4 PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see "Configuring Gigabit Ethernet Policers" on page 539, "Configuring MAC Address Filtering" on page 544, and "Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview" on page 559.

NOTE: When using the **show interfaces extensive** command with a 100-Gigabit Ethernet Type 4 PIC with CFP, the "Filter statistics" section will not be displayed because the hardware does not include those counters.

Release History Table

Release	Description
13.2	Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC.

Related Documentation

- 100-Gigabit Ethernet Type 4 PIC with CFP Overview on page 441
- Configuring Gigabit Ethernet Policers on page 539
- Configuring MAC Address Filtering on page 544
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

In Junos OS Release 10.4 and later, you can configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with routers using 100 gigabit Ethernet interfaces from other vendors by using the **forwarding-mode** statement with the **vlan-steering** option at the **[edit chassis fpc *slot* pic *slot*]** hierarchy level. On ingress, the router compares the outer VLAN ID against the user-defined VLAN ID and VLAN mask combination and steers the packet accordingly. You can program a custom VLAN ID and corresponding mask for PFE0.

General information on the VLAN steering mode:

- In VLAN steering mode, the SA multicast parameters are not used for packet steering.
- In SA multicast bit steering mode, the VLAN ID and VLAN masks are not used for packet steering.

- Configuration to set the packet distribution mode and VLAN steering rule is done through CLI commands. Both CLI commands result in a PIC reboot.
- There are three possible tag types of ingress packet:
  - Untagged ingress packet—The packet is sent to PFE1.
  - Ingress packet with one VLAN—The packet is forwarded to the corresponding PFE based on the VLAN ID.
  - Ingress packet with two VLANs—The packet is forwarded to the corresponding PFE based on the outer VLAN ID.
- If no VLAN rule is configured, all tagged packets are distributed to PFE0.
- VLAN rules describe how the router distributes packets. Two VLAN rules are provided by the CLI:
  - Odd-Even rule—Odd number VLAN IDs go to PFE1; even number of VLAN IDs go to PFE0.
  - Hi-Low rule—VLAN IDs 1 through 2047 go to PFE0; VLAN IDs 2048 through 4096 go to PFE1.
- When the 100-Gigabit Ethernet Type 4 PIC with CFP is configured in VLAN steering mode, it can be configured in a two physical interfaces mode or in aggregate Ethernet (AE) mode:
  - Two physical interfaces mode—When the PIC is in the two physical interfaces mode, it creates the physical interfaces **et-*x/0/0*:*0*** and **et-*x/0/0*:*1***. Each physical interface can configure its own logical interface and VLAN. The CLI enforces the following restrictions at the commit time:
    - The VLAN ID configuration must comply with the selected VLAN rule.
    - The previous restriction implies that the same VLAN ID cannot be configured on both physical interfaces.
  - AE mode—When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one AE physical interface. The PIC egress traffic is based on an AE internal hash algorithm. The PIC ingress traffic steering is based on the customized VLAN ID rule. The CLI enforces the following restrictions at the commit time:
    - The PICs AE working in VLAN steering mode includes both links of that PIC, and only the links of that PIC.
    - The PIC AE working in SA multicast steering mode can include more than one 100-Gigabit Ethernet Type 4 PIC with CFP to achieve more than 100 gigabit Ethernet capacity.

To configure SA multicast mode, use the **set chassis fpc *slot* pic *slot* forwarding-mode sa-multicast** command.

SA Multicast Mode

To configure SA multicast mode on a Juniper Networks 100-Gigabit Ethernet Type 4 PIC with CFP in FPC 0, PIC 0 for interconnection with another Juniper Networks 100-Gigabit Ethernet PIC, use the **set chassis fpc *slot* pic *slot* forwarding-mode sa-multicast** command. You can use the **show forwarding-mode** command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
forwarding-mode {
  sa-multicast;
}
```

VLAN Steering Mode

To configure the Juniper Networks 100-Gigabit Ethernet Type 4 PIC with CFP for VLAN steering mode for interoperation with a 100 gigabit Ethernet interface from another vendor's router, use the **set chassis fpc *slot* pic *slot* forwarding-mode vlan-steering** command with the **vlan-rule (high-low | odd-even)** statement. You can use the **show forwarding-mode** command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
forwarding-mode {
  vlan-steering {
    vlan-rule odd-even;
  }
}
```

Related Documentation

- forwarding-mode (100-Gigabit Ethernet) on page 1174
- sa-multicast (100-Gigabit Ethernet) on page 1350
- vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1443
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1444

100-Gigabit Ethernet Type 5 PIC with CFP Overview

The 100-Gigabit Ethernet PIC is a 1-port 100-Gigabit Ethernet Type 5 PIC with C form-factor pluggable transceiver (CFP) with model number PF-1CGE-CFP.

The following features are supported on 100-Gigabit Ethernet Type 5 PIC with CFP:

- Access to all 100-Gigabit Ethernet port counters through SNMP.
- Logical interface–level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Channels defined by two stacked VLAN tags.
- Channels defined by **flex-vlan-tagging**.
- IP service for stacked VLAN tags.
- Defining the rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.

NOTE: Only the Tag Protocol Identifier (TPID) 0x8100 is supported.

- Interface encapsulations, such as the following:
  - **untagged**—Default encapsulation, when other encapsulation is not configured.
    - You can configure only one logical interface (unit 0) on the port.
    - You cannot include the **vlan-id** statement in the configuration of the logical interface.
  - **vlan-tagging**—Enable VLAN tagging for all logical interfaces on the physical interface.
  - **stacked-vlan-tagging**—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
  - **ethernet-ccc**—Ethernet cross-connect.
  - **ethernet-tcc**—Ethernet translational cross-connect.
  - **vlan-ccc**—802.1Q tagging for a cross-connect.
  - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect.
  - **extended-vlan-ccc**—Standard TPID tagging for an Ethernet cross-connect.
  - **extended-vlan-tcc**—Standard TPID tagging for an Ethernet translational cross-connect.
  - **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration.
  - **ethernet-vpls**—Ethernet virtual private LAN service.
  - **vlan-vpls**—VLAN virtual private LAN service.
- The following Layer 3 protocols are also supported:
  - IPv4
  - IPv6
  - MPLS
- CFP Multi-Source Agreement (MSA) compliant Management Data Input/Output (MDIO) control features (transceiver dependent).
- 802.3 link aggregation:
  - The configuration of the 100-Gigabit Ethernet Type 5 PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
- Interoperability mode—Interoperability with the 100-Gigabit Ethernet Type 4 PIC with CFP through configuration in **sa-multicast** forwarding mode.
- Juniper Networks enterprise-specific Ethernet Media Access Control (MAC) MIB

- The 100-Gigabit Ethernet Type 5 PIC with CFP supports all Gigabit Ethernet PIC classification, firewall filters, queuing model, and Layer 2 rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see “Configuring Gigabit Ethernet Policers” on page 539, “Configuring MAC Address Filtering” on page 544, and “Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview” on page 559.
- A Type 5 FPC can support up to two 100-Gigabit Ethernet PICs. Both the PICs (that is, PIC 0 and PIC 1) can be offline or online independently.

The following features are not supported on the 100-Gigabit Ethernet Type 5 PIC with CFP:

- MAC filtering, accounting, and policing for destination MAC at the logical interface level.

NOTE: **Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This configuration enables the OSPF protocol to work.**

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.
- Multiple TPIDs.
- IP service for nonstandard TPID.

Table 38 on page 452 lists the capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP.

Table 38: Capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP

Capability	Support
Maximum logical interfaces per PIC	4093
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.

Related Documentation

- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP on page 444
- Configuring Gigabit Ethernet Policers on page 539
- Configuring MAC Address Filtering on page 544
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559

100-Gigabit Ethernet Interfaces Interoperability

Juniper Networks Junos operating system (Junos OS) supports a variety of 100-Gigabit Ethernet interfaces. The 100-Gigabit Ethernet standard, introduced by IEEE 802.3ba-2010, enables transmission of Ethernet frames at the rate of 100 gigabits per second (Gbps).

It is used for very high speed transmission of voice and data signals across the numerous world-wide fiber-optic networks.

Interface interoperability refers to the ability of an interface to interoperate with other router interfaces. You can enable interoperability between different 100-Gigabit Ethernet interfaces by performing specific configuration tasks. The following sections list the 100-Gigabit Ethernet interfaces, corresponding interoperable interfaces, and links to the interoperability tasks and reference information.

Interoperability of the MIC-3D-1X100GE-CFP MIC with PICs on Other Routers

Table 39 on page 453 lists the Interoperability with the 100-Gigabit Ethernet MIC with CFP.

Table 39: 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP) Interoperability

Interoperates with...		For More Information...
T Series	100-Gigabit Ethernet PIC with CFP (Type 4) (PD- 1CE-CFP-FPC4)	*Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode*

Interoperability of the MPC4E-3D-2CGE-8XGE MPC with PICs on Other Routers

Table 40 on page 453 lists the Interoperability with the MPC4E.

Table 40: MPC4E Interoperability

Interoperates with...		For More Information...
T Series	100-Gigabit Ethernet PIC with CFP (Type 4) (PD-1CECFP- FPC4)	*Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode*

Interoperability of the P1-PTX-2-100GE-CFP PIC with PICs on Other Routers

Table 41 on page 453 lists the Interoperability with 100-Gigabit Ethernet PIC with CFP (Type 5).

Table 41: 100-Gigabit Ethernet PIC with CFP (Type 5) (P1-PTX-2-100GE-CFP) Interoperability

Interoperates with...	For More Information...

Table 41: 100-Gigabit Ethernet PIC with CFP (Type 5) (P1-PTX-2-100GE-CFP) Interoperability (continued)

T Series	100-Gigabit Ethernet PIC with CFP (Type 4) (PD- 1CE-CFP-FPC4)	"Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP" on page 456 "Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4" on page 460

Interoperability of the PD-1CE-CFP-FPC4 PIC with PICs or MICs on Other Routers

Table 42 on page 454 lists the 100-Gigabit Ethernet PIC with CFP (Type 4).

Table 42: 100-Gigabit Ethernet PIC with CFP (Type 4) PD-1CE-CFP-FPC4 Interoperability

Interoperates with...		For More Information...
T Series	100-Gigabit Ethernet PIC with CFP (Type 5) (PF- 1CGE-CFP)	"Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4" on page 457 **forwarding-mode** **sa-multicast**
MX Series	100-Gigabit Ethernet MIC with CFP (MIC3-3D- 1X100GE-CFP)	*Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode*
	100-Gigabit Ethernet ports on the MPC4E	*Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode*
PTX Series	100-Gigabit Ethernet PIC with CFP (Type 5) (P1- PTX-2-100GE-CFP)	"Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP" on page 456 "Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4" on page 460

Related Documentation

- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 457
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 460
- *Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode*
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode*

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP by:

- Enabling source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP.
- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the 100-Gigabit Ethernet PIC PF-1CGE-CFP. The 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 looks at the bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. When the PIC sends out a packet, the multicast bit is set based on the egress Packet Forwarding Engine number (0 or 1).

The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.

PD-1CE-CFP-FPC4 uses two 50 Gpbs Packet Forwarding Engines to achieve 100 Gbps throughput. The 50-Gigabit Ethernet physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both the physical interfaces. You must configure the physical interfaces on PD-1CE-CFP-FPC4 in static link aggregation group (LAG) mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to the 100-Gigabit Ethernet PIC PF-1CGE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

NOTE: If you try to enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP without configuring PD-1CE-CFP-FPC4 (with two 50-Gigabit Ethernet interfaces) in static LAG mode, then there are issues in forwarding or routing protocols. For example, if you create two untagged logical interfaces—one each on the two 50-Gigabit Ethernet interfaces—on PD-1CE-CFP-FPC4 and one untagged logical interface on PF-1CGE-CFP, then PF-1CGE-CFP does not learn about one of the 50-Gigabit Ethernet interfaces on PD-1CE-CFP-FPC4.

Related Documentation

- **forwarding-mode on page 1174**
- **sa-multicast on page 1350**
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 457

- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP

You can enable interoperability between the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 and the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP by:

- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.
- Configuring source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP.

SA multicast bit steering mode uses the multicast bit in the source MAC address for packet steering.

NOTE: **When SA multicast bit steering mode is configured on a PTX Series Packet Transport Router 100-Gigabit Ethernet port, VLANs are not supported for that port.**

The 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 uses two 50-Gpbs Packet Forwarding Engines to achieve 100-Gbps throughput. The 50-Gigabit Ethernet physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both the physical interfaces. You must configure the physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 in static link aggregation group (LAG) mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP.

On the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4, ingress packets are forwarded to either Packet Forwarding Engine number 0 or 1 based on the SA multicast bit in the received packet. The SA multicast bit of egress packets is set based on whether the packet is forwarded from Packet Forwarding Engine number 0 or 1. As the default packet steering mode is SA multicast bit steering mode, no configuration is necessary to enable this mode.

On the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP, the SA multicast bit is ignored in ingress packets. When SA multicast bit steering mode is enabled, the SA multicast bit in the egress packets is set to 0 or 1 based on the flow hash value that is computed internally by the Packet Forwarding Engine complex for each packet. No CLI configuration is required to generate the flow hash value as this computation is done automatically. The flow hash algorithm uses fields in the packet header to compute the flow hash value. By default, the SA multicast bit is set to 0 in egress packets. You must configure SA multicast bit steering mode to enable interoperability with the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 .

NOTE: If you try to enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP without configuring PD-1CE-CFP-FPC4 (with two 50-Gigabit Ethernet interfaces) in static LAG mode, then there are issues in forwarding or routing protocols. For example, if you create two untagged logical interfaces—one each on the two 50-Gigabit Ethernet interfaces—on the PD-1CE-CFP-FPC4 and one untagged logical interface on the P1-PTX-2-100GE-CFP, then P1-PTX-2-100GE-CFP does not learn about one of the 50-Gigabit Ethernet interfaces on PD-1CE-CFP-FPC4.

Related Documentation

- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 460
- **sa-multicast on page 1351**
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP by performing the following tasks:

- Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP on page 457
- Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface on page 458

Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP, you need to enable source address (SA) multicast bit steering mode on PF-1CGE-CFP.

To configure SA multicast mode on PF-1CGE-CFP:

1. Specify the FPC and PIC information on the chassis.

   [edit]
   user@host# **edit chassis fpc** ***slot*** **pic** ***slot***

   For example:

   [edit]
   user@host# **edit chassis fpc 1 pic 0**

2. Configure the interoperation mode (SA multicast bit steering mode).

   [edit chassis fpc *slot* pic *slot*]

user@host# **set forwarding-mode sa-multicast**

For example:

[edit fpc 1 pic 0]
user@host# **set forwarding-mode sa-multicast**

3. Verify the configuration.

[edit]
user@host# **show chassis**

```
fpc 1 {
    pic 0 {
        forwarding-mode {
            sa-multicast;
        }
    }
}
```

NOTE: The default packet steering mode for the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.

See Also

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454
- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP or P1-PTX-2-100GE-CFP, you need to configure the two 50-Gigabit Ethernet physical interfaces on PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface. This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to PF-1CGE-CFP or P1-PTX-2-100GE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-*fpc/pic/*0:0 and et-*fpc/pic/*0:1, where *fpc* is the FPC slot number and *pic* is the PIC slot number. For example, to configure two physical interfaces for PIC slot 0 in FPC slot 5:

1. Specify the number of aggregated Ethernet interfaces to be created.

[edit chassis]
user@host# **set aggregated devices ethernet device-count *count***

For example:

[edit chassis]
user@host# **set aggregated devices ethernet device-count 1**

2. Specify the members to be included within the aggregated Ethernet bundle.

   [edit interfaces]
   user@host# **set *interface-name* gigether-options 802.3ad *bundle***

   The following example shows how to configure two physical interfaces for PIC 0 on a T1600 router.

   [edit interfaces]
   user@host# **set et-5/0/0:0 gigether-options 802.3ad ae0**
   user@host# **set et-5/0/0:1 gigether-options 802.3ad ae0**

3. Verify the configuration at the chassis.

   [edit]
   user@host# **show chassis**

```
aggregated-devices {
    ethernet {
        device-count 1;
    }
}
```

4. Verify the configuration at the interface.

   [edit]
   user@host# **show interfaces**

```
et-5/0/0:0 {
    gigether-options {
        802.3ad ae0;
    }
}
et-5/0/0:1 {
    gigether-options {
        802.3ad ae0;
    }
}
```

See Also

- Configuring Junos OS for Supporting Aggregated Devices on page 129
- **802.3ad on page 1063**

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454

- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP, you must enable source address (SA) multicast bit steering mode on P1-PTX-2-100GE-CFP.

NOTE: When you configure the SA multicast bit steering mode on the PTX Series PIC P1-PTX-2-100GE-CFP, we recommend that you do not configure the PIC ports as member links of an aggregated Ethernet interface because this prevents load balancing on the peering T Series PIC PD-1CE-CFP-FPC4. This T Series PIC must be in aggregated Ethernet mode to share bandwidth between its two 50-Gigabit Ethernet interfaces.

To configure SA multicast bit steering mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP:

1. Specify the FPC, PIC, and port information on the chassis.

   [edit]
   user@host# **edit chassis fpc *slot* pic *slot* port *port-number***

   For example:

   [edit]
   user@host# **edit chassis fpc 1 pic 0 port 0**

2. Configure the interoperation mode (SA multicast bit steering mode).

   [edit chassis fpc 1 pic 0]
   user@host# **set forwarding-mode sa-multicast**

3. Verify the configuration.

   [edit]
   user@host# **show chassis**

```
fpc 1 {
    pic 0 {
        port 0 {
            forwarding-mode {
                sa-multicast;
```



```
                }
            }
        }
    }
```

NOTE: As the default packet steering mode for the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 is SA multicast bit steering mode, no configuration is necessary to enable this mode.

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP or P1-PTX-2-100GE-CFP, you need to configure the two 50-Gigabit Ethernet physical interfaces on PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface. This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to PF-1CGE-CFP or P1-PTX-2-100GE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-*fpc*/*pic*/0:0 and et-*fpc*/*pic*/0:1, where *fpc* is the FPC slot number and *pic* is the PIC slot number. For example, to configure two physical interfaces for PIC slot 0 in FPC slot 5:

1. Specify the number of aggregated Ethernet interfaces to be created.

   [edit chassis]
   user@host# **set aggregated devices ethernet device-count *count***

   For example:

   [edit chassis]
   user@host# **set aggregated devices ethernet device-count 1**

2. Specify the members to be included within the aggregated Ethernet bundle.

   [edit interfaces]
   user@host# **set *interface-name* gigether-options 802.3ad *bundle***

   The following example shows how to configure two physical interfaces for PIC 0 on a T1600 router.

   [edit interfaces]
   user@host# **set et-5/0/0:0 gigether-options 802.3ad ae0**
   user@host# **set et-5/0/0:1 gigether-options 802.3ad ae0**

3. Verify the configuration at the chassis.

   [edit]
   user@host# **show chassis**

   ```
   aggregated-devices {
       ethernet {
   ```

```
            device-count 1;
        }
    }
```

4. Verify the configuration at the interface.

   [edit]
   user@host# **show interfaces**
   ```
   et-5/0/0:0 {
       gigether-options {
           802.3ad ae0;
       }
   }
   et-5/0/0:1 {
       gigether-options {
           802.3ad ae0;
       }
   }
   ```

See Also

- Configuring Junos OS for Supporting Aggregated Devices on page 129
- **802.3ad on page 1063**

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP on page 456
- **sa-multicast on page 1351**
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454

CHAPTER 27

Configuring Gigabit Ethernet OTN Options

Gigabit Ethernet OTN Options

The following example shows the configuration settings for Gigabit Ethernet OTN options:

```
[edit interfaces ge-fpc/pic/port]
otn-options {
  bytes (otn-options) transmit-payload-type value;
  fec (efec | gfec | gfec-sdfec | none );
  (is-ma | no-is-ma);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  (local-loopback | no-local-loopback);
  (odu-ttim-action-enable | no-odu-ttim-action-enable);
  (otu-ttim-action-enable | no-otu-ttim-action-enable);
  odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    number-of-frames value;
    (no-start-measurement | start-measurement;
  }
  odu-signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  (prbs | no-prbs);
  preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
    odu-backward-frr-enable | no-odu-backward-frr-enable;
    odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable;
  }
  rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    oc192;
    otu4;
    (pass-through | no-pass-through);
  }
  signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs
    number);
  transport-monitoring;
  trigger trigger-identifier;
  tti tti-identifier;
}
```

NOTE: The Gigabit Ethernet interface and the XENPAK interface support the read/write overhead bytes only for the APS/PPC (bytes 0 through 3).

You can use the following show commands to view the OTN configuration:

- **show interfaces extensive**—See the CLI Explorer for command details.
- **show chassis hardware**—See the CLI Explorer for command details.

- **show chassis pic**—See the CLI Explorer for command details.

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- *Ethernet Interfaces Feature Guide for Routing Devices*

10-Gigabit Ethernet OTN Options Configuration Overview

MX240, MX480, MX960, MX2010, MX2020, T320, T640, T1600, PTX3000, and PTX5000 routers support Optical Transport Network (OTN) interfaces, including the 10-Gigabit Ethernet DWDM OTN PIC, and provide ITU-T G.709 support. Use the **set** otn-options statement at the **[edit interfaces if-*fpc/pic/port*]** hierarchy level to configure the OTN options.

MX2020, MX2010, MX960, MX480, and MX240 routers support OTN interfaces on MPC5E and MPC6E. MPC5E-40G10G and MPC5EQ-10G40G support OTN on 10-Gigabit Ethernet interfaces but not on 40-Gigabit Ethernet interfaces. The OTN MIC MIC6-10G-OTN on MPC6E supports OTN on 10-Gigabit Ethernet interfaces on MX2020 and MX2010 routers. OTN support on the specified MX Series routers includes:

- International Telecommunications Union (ITU)-standard OTN performance monitoring and alarm management
- Transparent transport of 24 10-Gigabit Ethernet signals with optical channel data unit 2 (ODU2) and ODU2e framing on a per-port basis
- Pre-forward error correction (pre-FEC)-based bit error rate (BER). Fast reroute (FRR) uses the pre-FEC BER as an indication of the condition of an OTN link.

To configure the OTN options on the specified MX routers, use the **set** otn-options statement at the **[edit interfaces *interfaceType-fpc/pic/port*]** hierarchy level.

Related Documentation

- otn-options on page 1289
- *Ethernet Interfaces Feature Guide for Routing Devices*

Ethernet DWDM Interface Wavelength Overview

Dense wavelength-division multiplexing (DWDM) interfaces are supported on 10-Gigabit Ethernet DWDM PICs, MICs, and MPCs; the 10-Gigabit Ethernet LAN/WAN OTN PIC; and the 100-Gigabit Ethernet DWDM OTN PIC. When a tunable optic transceiver is available, you can configure the DWDM interfaces with full C-band International Telecommunication Union (ITU)-Grid tunable optics, as defined in the following specifications:

- *Intel TXN13600 Optical Transceiver I2C Interface and Customer EEPROM Preliminary Specification*, July 2004.
- *I2C Reference Document for 300-Pin MSA 10G and 40G Transponder*, Edition 4, August 04, 2003.

By default, the wavelength is 1550.12 nanometers (nm), which corresponds to 193.40 terahertz (THz).

Related Documentation

- Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength on page 466
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **wavelength on page 985**

Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength

To configure the wavelength on 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) and OTN interfaces, include the **wavelength** statement at the **[edit interfaces *interface-name* optics-options]** hierarchy level:

```
[edit interfaces interface-name optics-options]
wavelength nm;
```

To display the currently tuned wavelength and frequency for the interface, use the **show interfaces *interface-name*** operational mode command.

For interface diagnostics, issue the **show interfaces diagnostics optics *interface-name*** operational mode command.

Table 43 on page 466 shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

Table 43: Wavelength-to-Frequency Conversion Matrix

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
1528.38	196.15	1542.14	194.40	1556.15	192.65
1528.77	196.10	1542.54	194.35	1556.55	192.60
1529.16	196.05	1542.94	194.30	1556.96	192.55
1529.55	196.00	1543.33	194.25	1557.36	192.50
1529.94	195.95	1543.73	194.20	1557.77	192.45
1530.33	195.90	1544.13	194.15	1558.17	192.40
1530.72	195.85	1544.53	194.10	1558.58	192.35
1531.12	195.80	1544.92	194.05	1558.98	192.30
1531.51	195.75	1545.32	194.00	1559.39	192.25
1531.90	195.70	1545.72	193.95	1559.79	192.20

Table 43: Wavelength-to-Frequency Conversion Matrix (continued)

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
1532.29	195.65	1546.12	193.90	1560.20	192.15
1532.68	195.60	1546.52	193.85	1560.61	192.10
1533.07	195.55	1546.92	193.80	1561.01	192.05
1533.47	195.50	1547.32	193.75	1561.42	192.00
1533.86	195.45	1547.72	193.70	1561.83	191.95
1534.25	195.40	1548.11	193.65	1562.23	191.90
1534.64	195.35	1548.51	193.60	1562.64	191.85
1535.04	195.30	1548.91	193.55	1563.05	191.80
1535.43	195.25	1549.32	193.50	1563.45	191.75
1535.82	195.20	1549.72	193.45	1563.86	191.70
1536.22	195.15	1550.12	193.40	1564.27	191.65
1536.61	195.10	1550.52	193.35	1564.68	191.60
1537.00	195.05	1550.92	193.30	1565.09	191.55
1537.40	195.00	1551.32	193.25	1565.50	191.50
1537.79	194.95	1551.72	193.20	1565.90	191.45
1538.19	194.90	1552.12	193.15	1566.31	191.40
1538.58	194.85	1552.52	193.10	1566.72	191.35
1538.98	194.80	1552.93	193.05	1567.13	191.30
1539.37	194.75	1553.33	193.00	1567.54	191.25
1539.77	194.70	1553.73	192.95	1567.95	191.20
1540.16	194.65	1554.13	192.90	1568.36	191.15
1540.56	194.60	1554.54	192.85	1568.77	191.10
1540.95	194.55	1554.94	192.80		

Table 43: Wavelength-to-Frequency Conversion Matrix (continued)

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
1541.35	194.50	1555.34	192.75		
1541.75	194.45	1555.75	192.70		

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 465
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **wavelength on page 985**

Understanding the P1-PTX-24-10G-W-SFPP PIC

Starting from Junos OS Release 14.2, a 24–port 10-Gigabit Ethernet OTN PIC—P1-PTX-24-10G-W-SFPP—is supported on the FPC-PTX-P1-A and FPC2-PTX-P1A FPCs in PTX5000 routers, and the FPC-SFF-PTX-P1-A and FPC-SFF-PTX-T FPCs in PTX3000 routers. The P1-PTX-24-10G-W-SFPP PIC provides twenty-four 10-Gigabit Ethernet interfaces, that are independently configurable in LAN PHY or WAN PHY framing mode or in optical channel transport unit in OTU2e, OTU1e, or OTU2 mode.

The following sections explain this PIC in detail:

- Interface Features on page 468
- Layer 2 and Layer 3 Features on page 470
- OTN Alarms and Defects on page 471
- TCA Alarms on page 472

Interface Features

The following interface features are supported on the P1-PTX-24-10G-W-SFPP PIC:

- Twenty-four 10-Gigabit Ethernet interfaces, which are independently configurable in LAN PHY or WAN PHY mode or in OTU2e, OTU1e, or OTU2 signal mode. Each interface is terminated by means of a CFP2 transceiver.
- The interfaces are named with prefix *et*.
- Gigabit Ethernet local loopback.
- Link-level pause frames—You can halt the Ethernet interface from transmitting packets for a configured period of time.
- Interface hold timer and interface damping—You can set the **hold-time** statement (in milliseconds) to damp interface transitions.
- External clock.
- Nonstandard tag protocol identifier (TPID):

  - For each 10-Gigabit Ethernet port, you can configure up to eight TPIDs by using the **tag-protocol-id** statement at the **[edit interfaces interface-name gigether-options ethernet-switch-profile]** hierarchy level.
  - The **tag-protocol-id** statement can be configured only on the first port (port 0) of the PIC. If any other (nonzero) port has the **tag-protocol-id** configuration, the Routing Engine registers an error in the system log and the configuration is ignored.
  - The **tag-protocol-id** statement configured on port 0 of the PIC also applies to the rest of the ports on that PIC.
- Generic forward error correction (GFEC),ultra forward error correction (UFEC), enhanced forward error correction (EFEC), and no-FEC modes of operation are supported.
- Diagnostics tools:
  - Line loopback
  - Local loopback
- Fast reroute (FRR)—Based on configurable pre-FEC, bit error rate (BER) is supported and is configured using the **ber-threshold-signal-degrade** statement at the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level.
- *jnx-ifotn.mib* and *otn-mib* as defined in RFC 3591. Note that according to Junos OS security standard, configurable parameters are not supported through SNMP. Only the *get* operation is available through SNMP.
- FEC statistics—corrected errors and corrected error ratio.
- OTN payload pseudorandom binary sequence (PRBS) generation and checking by enabling or disabling PRBS with the **prbs** or **no-prbs** statement at the **[edit interfaces *interface-name* otn-options]** hierarchy level.
- At the physical interface level, **flexible-ethernet-service**, **ethernet-ccc**, and **ethernet-tcc** encapsulations are supported. For **flexible-ethernet-service** encapsulation, the logical level supports **enet2**, **vlan-ccc**, and **vlan-tcc** encapsulations.
- At the logical interface level **dix**, **vlan-ccc**, and **vlan-tc**c encapsulations are supported.
- SNMP management of the PIC based on RFC 3591, Definitions of Managed Objects for the Optical Interface Type:
  - Set functionality
  - Juniper Networks Black-Link MIB
  - IFOTN MIB
  - Optics MIB
  - FRU MIB
- 15-minute and 1-day performance monitoring and historic statistics.
  - Near-end and far-end performance monitoring
  - Threshold-crossing alerts

  - BER performance monitoring
  - FEC performance monitoring
  - Optical performance monitoring

The following features are not supported on the P1-PTX-24-10G-W-SFPP PIC:

- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per queue flow control
- Generic framing procedure-framed (GFP-F) mapping modes over OTN
- General communication channel (GCC)
- OTN interface-level Automatic Protection Switching (APS)
- Insertion, monitoring, and display of OTN header overhead byte
- Optical harness support
- Transport interface and state model (GR-1093)
- Trace tone support

Layer 2 and Layer 3 Features

The following Layer 2 and Layer 3 features are supported on the P1-PTX-24-10G-W-SFPP PIC:

- MAC detect link up and link down based on local fault signal or remote fault signal.
- MAC statistics.
- Flow control.
- MAC oversized packet counters based on default MTU value or user-configured MTU value.
- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Per logical interface source address MAC accounting.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.

- Aggregated Ethernet supports 64 child links that can be configured using the **set chassis aggregated-devices maximum-links** configuration command.
- Maximum of 1024 logical interfaces on an aggregated Ethernet physical interface.
- Support for V LAN tagging, flexible VLAN tagging, and stacked VLAN tagging.
- LACP.
- Link protection.
- 802.3 ah OAM.
- 802.1 ag OAM.
- MPLS FRR.
- SNMP.
- Supports per-VLAN queuing (using Packet Forwarding Engine).

OTN Alarms and Defects

The following OTN alarms and defects are supported on the P1-PTX-24-10G-W-SFPP PIC:

- LOS—Loss Of Signal
- LOF—Loss Of Frame
- LOM—Loss Of Multiframe
- SSF—Server Signal Failure
- TSF—Trail Signal Fail
- OTU-FEC-DEG—Forward Error Correction Degraded
- OTU-FEC-EXE—Excessive Errors, FEC_FAIL from the transponder
- OTU-AIS—Alarm Indication Signal or all ones signal
- OTU-BDI—Backward Defect Identification
- OTU-IAE—Incoming Alignment Error
- OTU-BIAE—Backward Incoming Alignment Error
- OTU-TTIM—Destination Access Point Identifier [DAPI], Source Access Point Identifier [SAPI], or both mismatch from expected to received
- OTU-SD—Signal Degrade
- OTU-SF—Signal Fail
- CSF—Client Signal Failure
- ODU-LCK—(ODU lock triggers for PM [path monitoring])
- ODU-AIS—(alarm indication signal or all ones signal)
- ODU-OCI—(open connection error)

- ODU-BDI—(backward defect indication)
- ODU-IAE—(incoming alignment error)
- ODU-DAPI-TTIM—DAPI or DAPI/SAPI mismatch from expected to receive
- ODU-SAPI-TTIM—SAPI or DAPI/SAPI mismatch from expected to receive
- ODU-BEI—Backward Error Indication
- ODU-SSF—Server Signal Fail
- ODU-TSF—Trail Signal Fail
- ODU-SD—Signal Degrade
- ODU-SF—Signal Fail
- OPU-PTM—Payload Type Mismatch

TCA Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UAS)

Related Documentation

- Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC on page 472

Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC

Starting from Junos OS Release 14.2, a 24–port 10-Gigabit Ethernet OTN PIC—P1-PTX-24-10G-W-SFPP—is supported on the FPC-PTX-P1-A and FPC2-PTX-P1A FPCs in PTX5000 routers, and the FPC-SFF-PTX-P1-A and FPC-SFF-PTX-T FPCsin PTX3000 routers. To configure an OTN interface on the P1-PTX-24-10G-W-SFPP PIC, you must configure interface-specific options and the OTN-related options for the interface.

To configure the interface-specific options:

1. Go to the **[edit interface *interface-name*]** hierarchy level, where *interface-name* is in the et-*fpc/pic/port* format.

   [edit]
   user@host# **edit interfaces *interface-name***

2. Configure the VLAN tagging option on the OTN interface to enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

   [edit interfaces *interface-name*]
   user@host# **set vlan-tagging**

3. Configure the maximum transmission unit (MTU) size in bytes for the interface.

   [edit interfaces *interface-name*]
   user@host# **set mtu *bytes***

4. Configure a VLAN ID for the interface.

   [edit interfaces *interface-name*]
   user@host# **set vlan-id *number***

5. Configure the family for the interface.

   [edit interfaces *interface-name*]
   user@host# **set family *family-name***

6. Configure an IP address for the interface.

   [edit interfaces *interface-name*]
   user@host# **set address *address***

To configure the OTN-related options on the interface:

1. Go to the **[edit interface *interface-name* otn-options]** hierarchy level:

   [edit interfaces *interface-name*]
   user@host# **edit otn-options**

2. Enable the OTN mode as OTU2e, OTU1e, or OTU2 for the interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set rate fixed-stuff-bytes|no-fixed-stuff-bytes|oc192**

   NOTE: **fixed-stuff-bytes** is for OTU2e rate, **no-fixed-stuff-bytes** is for OTU1e rate and **oc192** is for OTU2 rate. OTU2e and OTU1e rates are applicable for LAN PHY framing mode. OTU2 is applicable for WAN PHY framing mode. Framing mode is to set through the **set interfaces framing** configuration statement.

3. Enable the laser on the OTN interface. The laser is disabled by default for all OTN interfaces.

   [edit interfaces *interface-name* otn-options]
   user@host# **set laser-enable**

4. Set an trail trace identifier for the source access point and for the destination access point for ODU and OTU on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi | otu-dapi | otu-expected-receive-dapi | otu-expected-receive-sapi | otu-sapi)** ***tti-identifier***

5. Ignore the trigger for the defect or set the hold time.

   Configure the hold time for the defect trigger as:

   - *up* with a value—Wait for the hold time delay before clearing the alarm when the defect is absent on the OTN interface.
   - *down* with a value—Wait for the hold time delay before raising the alarm when the defect occurs for the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set trigger (oc-lof | oc-lom | oc-los | oc-tsf | odu-ais | odu-bdi | odu-bei | odu-iae | odu-lck | odu-oci | odu-sd | odu-ttim |opu-ptim | otu-ais | otu-bdi | otu-fec-deg | otu-fec-exe | otu-iae | otu-sd | otu-ttim) (hold-time (down** ***value*** **| up** ***value*****) | ignore)**

6. Enable the threshold crossing alarms for the OTN interface along with the trigger for the defect.

   [edit interfaces *interface-name* otn-options]
   user@host# **set tca (odu-tca-bbe | odu-tca-es | odu-tca-ses | odu-tca-uas | otu-tca-bbe | otu-tca-es | otu-tca-ses | otu-tca-uas) (enable-tca | no-enable-tca | threshold)**

7. Set the OTN header bytes as a transmit payload type from 0 bytes through 255 bytes for the packets that are transmitted on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set bytes transmit-payload-type** ***value***

8. Configure the forward error correction (FEC) mode as Generic Forward Error Correction (GFEC), Enhanced Forward Error Correction (EFEC), Ultra Forward Error Correction (UFEC), or no-FEC (none) for the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set fec (gfec | ufec | efec | none)**

9. Enable a consequent action as listed in the ITU-T G.798 standard for ODU trail trace identifier mismatch (TTIM) on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set odu-ttim-action-enable**

10. Enable a consequent action as listed in the ITU-T G.798 standard for OTU trail trace identifier mismatch (TTIM) on the OTN interface.


```
[edit interfaces interface-name otn-options]
user@host# set otu-ttim-action-enable
```

11. Configure the threshold value for signal degradation when an alarm needs to be raised. Configure the threshold value after signal degradation when the alarm needs to be cleared. When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear *value*** statement, then BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

```
[edit interfaces interface-name otn-options signal-degrade]
user@host# set ber-threshold-signal-degrade value
user@host# set ber-threshold-clear value
user@host# set interval value
```

12. Enable the following actions for the **preemptive-fast-reroute** statement:

    - Backward FRR—Insert the local pre-FEC status into the transmitted OTN frames and monitor the received OTN frames for the pre-FEC status.

    ```
    [edit interfaces interface-name otn-options preemptive-fast-reroute]
    user@host# set backward-frr-enable
    ```

    - Monitoring of signal degradation of pre-FEC OTN frames.

    ```
    [edit interfaces interface-name otn-options preemptive-fast-reroute]
    user@host# set signal-degrade-monitor-enable
    ```

Related Documentation

- Understanding the P1-PTX-24-10G-W-SFPP PIC on page 468
- **optics-options on page 1288**
- **otn-options on page 1289**
- **signal-degrade on page 971**
- **preemptive-fast-reroute on page 968**

100-Gigabit Ethernet OTN Options Configuration Overview

PTX Series routers support optical transport network (OTN) interfaces, including the 100-Gigabit DWDM OTN PIC, which supports:

- Transparent transport of two 100-Gigabit Ethernet signals with Optical Channel Transport Unit 4 (OTU4) framing.
- International Telecommunications Union (ITU)-standard OTN performance monitoring (PM) and alarm management.

- Dual polarization quadrature phase shift keying (DP-QPSK) modulation and soft-decision forward error correction (SD-FEC) for long haul and metro applications.
- Pre-forward error correction (pre-FEC)-based bit error rate (BER) monitoring. Pre-FEC BER monitoring uses the pre-FEC BER as an indication of the condition of an OTN link. See “Understanding Pre-FEC BER Monitoring and BER Thresholds” on page 477 for more information.

For more information about the 100-Gigabit DWDM OTN PIC, see *100-Gigabit DWDM OTN PIC* in the *PTX Series Interface Module Reference*.

PTX Series routers also support the 100-Gigabit Ethernet OTN PIC (P2-100GE-OTN), which provides four 100-Gigabit Ethernet interfaces, independently configurable in LAN PHY framing mode or in optical channel transport unit 4 (OTU4) mode. See “Understanding the P2-100GE-OTN PIC” on page 500 for more information.

See “Supported OTN Options on PTX Series Routers” on page 485 for a comparison of the features supported on PTX Series OTN PICs.

MX2020, MX2010, MX960, MX480, and MX240 routers support OTN interfaces on MPC5E and MPC6E. MPC5E-100G10G and MPC5EQ-100G10G support 100-Gigabit Ethernet OTN interfaces and 10-Gigabit Ethernet OTN interfaces on MX240, MX480, and MX960 routers. The OTN MIC MIC6-100G-CFP2 on MPC6E supports OTN on 100-Gigabit Ethernet interfaces on MX2020 and MX2010 routers. OTN support on the specified MX Series routers includes:

- International Telecommunications Union (ITU)-standard OTN performance monitoring (PM) and alarm management
- Transparent transport of two 100-Gigabit Ethernet signals with optical channel transport unit 4 (OTU4) framing.
- Generic forward error correction (Generic FEC)

To configure the OTN options for PTX Series routers and specific MX Series routers, use the **set** otn-options statement at the **[edit interfaces *interfaceType-fpc/pic/port*]** hierarchy level.

Use the **set** optics-options statement at the **[edit interfaces *interfaceType-fpc/pic/port*]** hierarchy level to configure the optics options.

Use the *show interfaces extensive*, **show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)**, and **show interfaces transport pm** commands to view optics and OTN PM information. To display the current time interval and clear the channel service unit (CSU) alarm and defect counters, use the **clear interfaces interval** command.

Related Documentation

- Configuring 100-Gigabit DWDM OTN PICs on page 482
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) on page 1803**

- optics-options on page 1288
- otn-options on page 1289

Understanding Pre-FEC BER Monitoring and BER Thresholds

Optical transport network (OTN) interfaces on PTX Series Packet Transport Routers support monitoring the condition of an OTN link by using the pre-forward error correction (pre-FEC) bit error rate (BER). The following PICs support pre-FEC BER monitoring:

- P1-PTX-2-100G-WDM
- P2-100GE-OTN
- P1-PTX-24-10G-W-SFPP

The PICs use forward error correction (FEC) to correct bit errors in the received data. As long as the pre-FEC BER is below the FEC limit, all bit errors are successfully identified and corrected and, therefore, no packet loss occurs. The system monitors the pre-FEC BER on each port. This gives an early warning of link degradation. By configuring an appropriate pre-FEC BER threshold and interval, you enable the PIC to take preemptive action before the FEC limit is reached. If this pre-FEC BER threshold logic is combined with MPLS fast reroute, then packet loss can be minimized or prevented.

You must specify both the signal degradation threshold (**ber-threshold-signal-degrade**) and the interval (**interval**) for the interface. The threshold defines the BER criteria for a signal degrade condition and the interval defines the minimum duration over which the BER must exceed the threshold before an alarm is raised. The relationship between the threshold and the interval is illustrated in Figure 32 on page 478. After an alarm is raised, if the BER returns to a level below the threshold clear value (**ber-threshold-clear**), the alarm is cleared.

Figure 32: Pre-FEC BER Monitoring

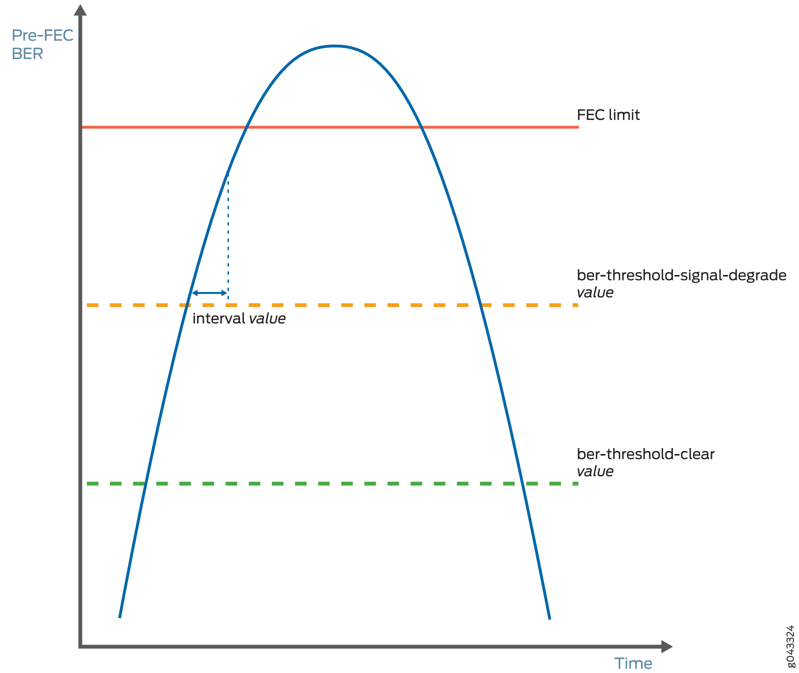

With pre-FEC BER monitoring enabled, when the configured pre-FEC BER signal degrade threshold is reached, the PIC stops forwarding packets to the remote interface and raises an interface alarm. Ingress packets continue to be processed. If pre-FEC BER monitoring is used with MPLS fast reroute or another link protection method, then traffic is rerouted to a different interface.

You can also configure backward fast reroute to insert the local pre-FEC status into transmitted OTN frames, notifying the remote interface of signal degradation. The remote interface can use the information to reroute traffic to a different interface. If you use pre-FEC BER monitoring together with backward fast reroute, then notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol.

Include the signal-degrade-monitor-enable and backward-frr-enable statements at the **[edit interfaces *interface-name* otn-options preemptive-fast-reroute]** hierarchy level to enable pre-FEC BER monitoring and backward fast reroute.

NOTE: When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the **signal-degrade-monitor-enable** and the **backward-frr-enable** statements.

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used.

When a received signal degrade alarm is active and backward fast reroute is enabled, a specific flag is inserted into the trasmitted OTN overhead. The remote PIC at the opposite end of the link monitors the OTN overhead, thus enabling both ends to initiate traffic rerouting in the event of a signal degrade condition. When the signal degrade condition is cleared, the OTN overhead flag is returned to a normal state.

The pre-FEC BER signal degrade threshold value defines a specific amount of system margin relative to the BER correction limit (or FEC limit) of the PIC's receive FEC decoder. Each PIC has a set FEC limit—it is intrinsic to the FEC decoder implementation.

NOTE: The examples below use Q^2-factor measurements (also known as Q-factor). Q^2-factor is expressed in units of decibels relative to a Q^2-factor of zero (dBQ). Q^2-factor enables you to describe system margin in linear terms in contrast to BER values, which are nonlinear in nature. After you determine the thresholds, you must convert the threshold values from Q^2-factor to BER to enter them in the CLI by using scientific notation. BER can be converted to Q^2-factor by using the following equation:

Q^2-factor = 20 * log10 (sqrt(2) * erfcinv(2 * *BER*))

TIP: To convert between Q^2-factor and BER in a spreadsheet program, you can approximate the values by using the following formulas:

- To calculate Q^2-factor:

  = 20 * LOG10(–NORMSINV(*BER*))

- To calculate BER:

  = 1 – NORMSDIST(10^(0.05 * *Q^2-factor*))

Table 44 on page 479 shows the relationship between the fixed FEC limit, the configurable signal degrade threshold, and the configurable clear threshold for different PICs. In this example, approximately 1 dBQ of system margin has been set between the FEC limit, signal degrade threshold, and clear threshold.

Table 44: Example—Signal Degrade and Clear Threshold Values at 1 dBQ

PIC	FEC Type	FEC Limit		Signal Degrade Threshold		Clear Threshold	
		Q^2-Factor	BER	Q^2-Factor	BER	Q^2-Factor	BER
P1-PTX-2-100G-WDM	SD-FEC	6.7 dBQ	1.5E–2	7.7 dBQ	7.5E–3	8.7 dBQ	3.0E–3
P2-100GE-OTN	G.709 GFEC	11.5 dBQ	8.0E–5	12.5 dBQ	1.1E–5	13.5 dBQ	1.0E–6

Table 44: Example—Signal Degrade and Clear Threshold Values at 1 dBQ (continued)

PIC	FEC Type	FEC Limit		Signal Degrade Threshold		Clear Threshold	
		Q^2-Factor	BER	Q^2-Factor	BER	Q^2-Factor	BER
P1-PTX-24-10G-W-SFPP	G.975.1 I.4 (UFEC)	9.1 dBQ	2.2E–3	10.1 dBQ	6.9E–4	11.1 dBQ	1.6E–4
	G.975.1 I.7 (EFEC)	9.6 dBQ	1.3E–3	10.6 dBQ	3.6E–4	11.6 dBQ	7.5E–5
	G.709 GFEC	11.5 dBQ	8.0E–5	12.5 dBQ	1.1E–5	13.5 dBQ	1.0E–6

To adjust the signal degrade threshold, you must first decide on a new system margin target and then calculate the respective BER value (using the equation to convert from Q^2-factor to BER). Table 45 on page 480 shows the values if 3 dBQ of system margin relative to the FEC limit is required for the signal degrade threshold (while maintaining the clear threshold at 1 dBQ relative to the signal degrade threshold).

NOTE: The choice of system margin is subjective, as you might want to optimize your thresholds based on different link characteristics and fault tolerance and stability objectives. For guidance about configuring pre-FEC BER monitoring and BER thresholds, contact your Juniper Networks representative.

Table 45: Example—Signal Degrade and Clear Thresholds After Configuration

PIC	FEC Type	FEC Limit		Signal Degrade Threshold		Clear Threshold	
		Q^2-Factor	BER	Q^2-Factor	BER	Q^2-Factor	BER
P1-PTX-2-100G-WDM	SD-FEC	6.7 dBQ	1.5E–2	9.7 dBQ	1.1E–3	10.7 dBQ	2.9E–4
P2-100GE-OTN	G.709 GFEC	11.5 dBQ	8.0E–5	14.5 dBQ	4.9E–8	15.5 dBQ	1.1E–9
P1-PTX-24-10G-W-SFPP	G.975.1 I.4 (UFEC)	9.1 dBQ	2.2E–3	12.1 dBQ	2.8E–5	13.1 dBQ	3.1E–6
	G.975.1 I.7 (EFEC)	9.6 dBQ	1.3E–3	12.6 dBQ	1.1E–5	13.6 dBQ	9.1E–7
	G.709 GFEC	11.5 dBQ	8.0E–5	14.5 dBQ	4.8E–8	15.5 dBQ	1.1E–9

Include the ber-threshold-signal-degrade, ber-threshold-clear, and interval statements at the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level to configure the BER thresholds and time interval.

NOTE: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

Supported Forward Error Correction Modes on MX Series Routers

Table 46 on page 481 lists the FEC modes that are supported on MX Series routers at the **[edit interfaces *interface-name* otn-options]** hierarchy level. Note that the term *NA* denotes that the statement is not applicable for that particular line card:

Table 46: FEC modes Supported on MX Series Routers

Line Card	FEC Mode	Port Speed	Junos Version
MPC5E-40G10G	***(gfec | efec | none | ufec)***	10G	13.3
MPC5E-100G10G	***(gfec | efec | none | ufec)***	10G and 100G (GFEC only)	13.3
MIC6-10G-OTN	***(gfec | efec | none | ufec)***	10G	13.3
MIC6-100G-CFP2	***(gfec | none)***	100G (GFEC only)	13.3
MIC3-100G-DWDM	***gfec | hgfec | sdfec***	100G	15.1F5

Related Documentation

- fec on page 948
- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477

Supported Forward Error Correction Modes on PTX Series Routers

Table 47 on page 481 lists the FEC modes that are supported on PTX Series routers at the **[edit interfaces *interface-name* otn-options]** hierarchy level.

Table 47: FEC Modes Supported on PTX Series Routers

Line Card	FEC Mode	Port Speed	Junos Version
P1-PTX-24-10G-W-SFPP	***(gfec | efec | none | ufec)***	10G	12.1X48, 12.3, 13.2 (PTX5000) 13.2R2 (PTX3000)
P2-10G-40G-QSFPP	***(gfec | efec | none | ufec)***	10G	14.1R2 (PTX5000) 15.1F6 (PTX3000)

Table 47: FEC Modes Supported on PTX Series Routers (continued)

Line Card	FEC Mode	Port Speed	Junos Version
P2-100GE-OTN	***(gfec | none)***	100G (GFEC only)	14.1
P1-PTX-2-100G-WDM	***(gfec-sdfec)***	100G	13.2 (PTX5000) 13.3 (PTX3000)
PTX-5-100G-WDM	***gfec | sdfec***	100G	15.1F6

Related Documentation

- **fec on page 948**
- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477

Configuring 100-Gigabit DWDM OTN PICs

PTX Series routers support optical transport network (OTN) interfaces, including the 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM). See "100-Gigabit Ethernet OTN Options Configuration Overview" on page 475.

To configure the 100-Gigabit DWDM OTN PIC:

1. Configure the interface wavelength.

   [edit interfaces *interface-name* optics-options]
   user@host# **set wavelength** *nm*

   See **wavelength**.

   NOTE: See *100-Gigabit DWDM OTN PIC Integrated Transceiver Optical Interface Specifications* for a list of wavelengths supported by the P1-PTX-2-100G-WDM PIC.

2. Enable the laser.

   [edit interfaces *interface-name* otn-options]
   user@host# **set laser-enable**

3. (Optional) Set the tca.

   [edit interfaces *interface-name* otn-options]
   user@host# **set tca** ***tca-identifier*** (enable-tca | no-enable-tca) (threshold *number* | threshold-24hrs *number*)

   See **tca**.

4. (Optional) Set the trace identifiers.

   [edit interfaces *interface-name* otn-options]
   user@host# **set tti** ***tti-identifier tti-identifier-name***

See tti.

5. (Optional) Specify defect triggers.

   [edit interfaces *interface-name* otn-options]
   user@host# **set trigger *trigger-identifier***

   See trigger.

6. (Optional) Enable VLAN tagging. See *Enabling VLAN Tagging*.

7. (Optional) Set the media MTU. See *Configuring the Media MTU*.

8. (Optional) Set the unit VLAN ID, family **inet**, and IP address.

   [edit interfaces *interface-name*]
   user@host# **set vlan-id *number***
   user@host# **set family inet**
   user@host# **set address *address***

9. (Optional) Enable pre-FEC BER signal-degrade monitoring and backward fast reroute to monitor the pre-FEC BER status of the link and to insert the local pre-FEC status into transmitted OTN frames.

   [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
   user@host# **set signal-degrade-monitor-enable**
   user@host# **set backward-frr-enable**

See signal-degrade-monitor-enable and backward-frr-enable.

10. (Optional) Configure the bit error rate (BER) thresholds for signal degradation used for monitoring the pre-forward error correction (pre-FEC) status of the OTN link.

    a. Set the BER signal-degrade threshold.

        [edit interfaces *interface-name* otn-options signal-degrade]
        user@host# **set ber-threshold-signal-degrade** ***value***

    b. Set the BER threshold to clear signal-degrade alarms.

        [edit interfaces *interface-name* otn-options signal-degrade]
        user@host# **set ber-threshold-clear** ***value***

    c. Set the time interval for signal-degrade collection. After the BER threshold for signal-degrade is crossed for ten consecutive intervals, an alarm is raised. If the BER threshold for signal-degrade clear is crossed for ten consecutive intervals, the alarm is cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

        [edit interfaces *interface-name* otn-options signal-degrade]
        user@host# **set interval** ***value***

        NOTE: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

        See ber-threshold-signal-degrade, ber-threshold-clear, and interval.

    NOTE: See "Understanding Pre-FEC BER Monitoring and BER Thresholds" on page 477 for more information about pre-FEC BER monitoring and determining BER threshold settings.

Related Documentation

- optics-options on page 1288
- otn-options on page 1289
- signal-degrade on page 971
- preemptive-fast-reroute on page 968

Supported OTN Options on PTX Series Routers

Table 48 on page 485 lists the statements that are supported on 100-Gigabit Ethernet PICs on PTX Series routers at the **[edit interfaces *interface-name* otn-options]** hierarchy level. Note that the term *NA* denotes that the statement is not applicable for that particular component:

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
bytes (otn-options)	transmit-payload-type *value*	13.2/13.3	14.1R2 14.2	14.2
fec	*(efec | gfec | gfec-sdfec | none | ufec)*	13.2/13.3	14.1R2 14.2 (gfec, none)	14.2
insert-odu-lck	-	13.2/13.3	14.1R2 14.2	14.2
insert-odu-oci	-	13.2/13.3	14.1R2 14.2	14.2
is-ma | no-is-ma	-	13.2/13.3	NA	14.2
laser-enable | no-laser-enable	-	13.2/13.3	14.1R2 14.2	14.2
line-loopback | no-line-loopback	-	13.2/13.3	14.1R2 14.2	14.2
local-loopback | no-local-loopback	-	13.2/13.3	14.1R2 14.2	14.2

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
odu-delay-management	bypass | no-bypass	13.2/13.3	NA	NA
	monitor-end-point | no-monitor-end-point	13.2/13.3	NA	NA
	number-of-frames *value*	13.2/13.3	NA	NA
	no-start-measurement | start-measurement	13.2/13.3	NA	NA
odu-signal-degrade	ber-threshold-clear *value*	NA NA	14.1R2 14.2	NA
	ber-threshold-signal-degrade *value*	NA NA	14.1R2 14.2	NA
	interval *value*	NA	14.1R2 14.2	NA
odu-ttim-action-enable | no-odu-ttim-action-enable	-	13.2/13.3	14.1R2 14.2	14.2
otu-ttim-action-enable | no-otu-ttim-action-enable	-	13.2/13.3	14.1R2 14.2	14.2
prbs | no-prbs	-	13.2/13.3	14.1R2 14.2	14.2
preemptive-fast-reroute	backward-frr-enable | no-backward-frr-enable	13.2/13.3	14.1R2 14.2	14.2
	signal-degrade-monitor-enable | no-signal-degrade-monitor-enable	13.2/13.3	14.1R2 14.2	14.2
	odu-backward-frr-enable | no-odu-backward-frr-enable	NA	14.1R2 14.2	NA
	odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable	NA	14.1R2 14.2	NA

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
rate	fixed-stuff-bytes | no-fixed-stuff-bytes	13.2/13.3	NA	14.2
	oc192	13.2/13.3	NA	14.2
	otu4	13.2/13.3	14.1R2 14.2	NA
	pass-through | no-pass-through	13.2/13.3	NA	NA
signal-degrade	ber-threshold-clear *value*	13.2/13.3	14.1R2 14.2	14.2
	ber-threshold-signal-degrade *value*	13.2/13.3	14.1R2 14.2	14.2
	interval *value*	13.2/13.3	14.1R2 14.2	14.2

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
tca	**odu-tca-bbe (enable-tca | no-enable-tca | threshold | threshold-24hrs)**	13.2/13.3	14.2	14.2
	odu-tca-bbe-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	NA	14.2
	odu-tca-es (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	odu-tca-es-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	NA	14.2
	odu-tca-ses (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	odu-tca-ses-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	NA	14.2
	odu-tca-uas (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	otu-tca-bbe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
		13.2/13.3	NA	14.2

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
	otu-tca-bbe-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)			
	otu-tca-es (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	otu-tca-es-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	NA	14.2
	otu-tca-fec-ber (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	otu-tca-ses (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	otu-tca-ses-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	NA	14.2
	otu-tca-uas (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	14.2	14.2
	otu-tca-uas-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.2/13.3	NA	14.2
transport-monitoring	-	NA	NA	14.2

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
trigger *trigger-identifier*	oc-lof (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	oc-lom (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	oc-los (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	oc-tsf (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	oc-wavelength-lock (hold-time (down | up) | ignore)	13.2/13.3	NA	14.2
	odu-ais (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-bdi (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-bei (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-iae (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-lck (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-oci (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-sd (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	odu-ttim	13.2/13.3		14.2

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
			14.1R2 14.2	
	opu-ptim (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-ais (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-bdi (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-fec-deg (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-fec-exe (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-iae (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-sd (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2
	otu-ttim (hold-time (down | up) | ignore)	13.2/13.3	14.1R2 14.2	14.2

Table 48: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000)	P2-100GE-OTN (PTX5000)	P1-PTX-24-10G-W-SFPP (PTX5000)
tti *tti-identifier*	odu-dapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	odu-expected-receive-dapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	odu-expected-receive-sapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	odu-sapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	otu-dapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	otu-expected-receive-dapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	otu-expected-receive-sapi *identifier*	13.2/13.3	14.1R2 14.2	14.2
	otu-sapi *identifier*	13.2/13.3	14.1R2 14.2	14.2

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

Supported OTN Options on MX Series Routers

Table 49 on page 493 lists the statements that are supported on 100-Gigabit Ethernet MICs on MX Series routers at the **[edit interfaces *interface-name* otn-options]** hierarchy level. Note that the term *NA* denotes that the statement is not applicable for that particular component:

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
bytes	transmit-payload-type *value*	NA	15.1F5
fec	*(efec | gfec | gfec-sdfec | hgfec| none | dsfec| ufec)*	13.3R3 (ufec)	15.1F5 (gfec,hgfec,sdfec)
insert-odu-lck	-	13.3R3	15.1F5
insert-odu-oci	-	13.3R3	15.1F5
is-ma | no-is-ma	-	13.3R3	15.1F5
laser-enable | no-laser-enable	-	13.3R3	15.1F5
line-loopback | no-line-loopback	-	13.3R3	15.1F5
local-loopback | no-local-loopback	-	13.3R3	15.1F5
odu-delay-management	bypass | no-bypass	NA	15.1F5
	monitor-end-point | no-monitor-end-point	NA	15.1F5
	number-of-frames *value*	NA	15.1F5
	no-start-measurement | start-measurement	NA	15.1F5
signal-degrade	ber-threshold-clear *value*	13.3R3	15.1F5
	ber-threshold-signal-degrade *value*	13.3R3	15.1F5
	interval *value*	13.3R3	15.1F5
odu-ttim-action-enable | no-odu-ttim-action-enable	-	13.3R3	15.1F5

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
otu-ttim-action-enable | no-otu-ttim-action-enable	-	13.3R3	15.1F5
prbs | no-prbs	-	13.3R3	15.1F5
preemptive-fast-reroute	backward-frr-enable | no-backward-frr-enable	13.3R3	15.1F5
	signal-degrade-monitor-enable | no-signal-degrade-monitor-enable	13.3R3	15.1F5
	odu-backward-frr-enable | no-odu-backward-frr-enable	NA	15.1F5
	odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable	NA	15.1F5
rate	fixed-stuff-bytes | no-fixed-stuff-bytes	13.3R3	15.1F5
	oc192	13.3R3	15.1F5
	otu4	13.3R3	15.1F5
	pass-through	13.3R3 (pass-through)	15.1F5 (pass-through)
signal-degrade	ber-threshold-clear *value*	13.3R3	15.1F5
	ber-threshold-signal-degrade *value*	13.3R3	15.1F5
	interval *value*	13.3R3	15.1F5
	q-threshold-signal-degrade	NA	15.1F5
	q-threshold-signal-degrade-clear	NA	15.1F5

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
tca	odu-tca-bbe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	odu-tca-bbe-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	odu-tca-es (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	odu-tca-es-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	odu-tca-ses (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	odu-tca-ses-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	odu-tca-uas (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-bbe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-bbe-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-es (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-es-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-fec-ber (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
		13.3R3	15.1F5

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
	otu-tca-ses (enable-tca | no-enable-tca | threshold | threshold-24hrs)		
	otu-tca-ses-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-uas (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
	otu-tca-uas-fe (enable-tca | no-enable-tca | threshold | threshold-24hrs)	13.3R3	15.1F5
transport-monitoring	-	NA	15.1F5

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
trigger *trigger-identifier*	oc-lof (hold-time (down | up) | ignore)	13.3R3	15.1F5
	oc-lom (hold-time (down | up) | ignore)	13.3R3	15.1F5
	oc-los (hold-time (down | up) | ignore)	13.3R3	15.1F5
	oc-tsf (hold-time (down | up) | ignore)	NA	15.1F5
	oc-wavelength-lock (hold-time (down | up) | ignore)	NA	15.1F5
	odu-ais (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-bdi (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-bei (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-iae (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-lck (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-oci (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-sd (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-tca-es	13.3R3	NA
	odu-tca-ses	13.3R3	NA
	odu-tca-uas	13.3R3	NA
	odu-ttim	13.3R3	15.1F5
	opu-ptim (hold-time (down | up) | ignore)	13.3R3	15.1F5

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
	otu-ais (hold-time (down | up) | ignore)	13.3R3	15.1F5
	otu-bdi (hold-time (down | up) | ignore)	13.3R3	15.1F5
	otu-fec-deg (hold-time (down | up) | ignore)	NA	15.1F5
	otu-fec-exe (hold-time (down | up) | ignore)	NA	15.1F5
	otu-iae (hold-time (down | up) | ignore)	13.3R3	15.1F5
	otu-sd (hold-time (down | up) | ignore)	13.3R3	15.1F5
	odu-tca-es	13.3R3	NA
	odu-tca-ses	13.3R3	NA
	odu-tca-uas	13.3R3	NA
	otu-ttim (hold-time (down | up) | ignore)	13.3R3	15.1F5

Table 49: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020)
tti *tti-identifier*	odu-dapi *identifier*y	13.3R3	15.1F5
	odu-dapi-first-byte-nul | no-odu-dapi-first-byte-nul	NA	15.1F5
	odu-expected-receive-dapi*identifier*	13.3R3	15.1F5
	odu-expected-receive-dapi-first-byte-nul| no-odu-expected-receive-dapi-first-byte-nul	NA	15.1F5
	odu-expected-receive-sapi *identifier*	13.3R3	15.1F5
	odu-sapi *identifier*y	13.3R3	15.1F5
	odu-sapi-first-byte-nul | no-odu-sapi-first-byte-nul	NA	15.1F5
	otu-dapi *identifier*	13.3R3	15.1F5
	otu-dapi-first-byte-nul | no-otu-dapi-first-byte-nul	NA	15.1F5
	otu-expected-receive-dapi *identifier*	13.3R3	15.1F5
	otu-expected-receive-dapi-first-byte-nul| no-odu-expected-receive-dapi-first-byte-nul	NA	15.1F5
	otu-expected-receive-sapi *identifier*	13.3R3	15.1F5
	otu-expected-receive-sapi-first-byte-nul| no-odu-expected-receive-dapi-first-byte-nul	NA	15.1F5
	otu-sapi *identifier*	13.3R3	15.1F5
	otu-sapi-first-byte-nul	NA	15.1F5

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
-

Understanding the P2-100GE-OTN PIC

Starting with Junos OS Release 14.1R2 and 14.2, a 100-Gigabit Ethernet OTN PIC—P2-100GE-OTN—is supported on the FPC2-PTX-P1A FPC in PTX5000 routers. The P2-100GE-OTN PIC provides 4-port 100-Gigabit Ethernet interfaces, which are independently configurable in LAN PHY framing mode or in optical channel transport unit 4 (OTU4) mode. Each interface is terminated by means of a CFP2 transceiver. The FPC2-PTX-P1A FPC supports two P2-100GE-OTN PICs, in which each 100-Gigabit Ethernet port is mapped to a Packet Forwarding Engine in the FPC.

Starting from Junos OS Release 15.1, you can configure the interfaces on the P2-100GE-OTN PIC on PTX5000 routers, to be a part of the mixed rates and mixed mode aggregated Ethernet bundles.

For information about mixed rates, see "Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles" on page 117.

Starting from Junos OS Release 15.1, you can configure port-based pseudowire class of service (CoS) classification which includes Layer 3 IPv4, IPv6, and MPLS classification for interfaces with ethernet-ccc encapsulation.

The following sections explain this PIC in detail:

Interface Features

The following interface features are supported on a P2-100GE-OTN PIC:

- 4-port 100-Gigabit Ethernet interfaces, which are independently configurable in LAN PHY framing mode or in OTU4 signal mode. Each interface is terminated by means of a CFP2 transceiver.
- Each port maps to a single Packet Forwarding Engine in the FPC2-PTX-P1A FPC.
- The interfaces are named with prefix *et*.
- Gigabit Ethernet local loopback.
- Link-level pause frames—You can halt the Ethernet interface from transmitting packets for a configured period of time.
- Interface hold timer and interface damping—You can set the **hold-time** statement (in milliseconds) to damp interface transitions.
- External clock

- Nonstandard tag protocol identifier (TPID):
  - For each 100-Gigabit Ethernet port, you can configure up to eight TPIDs by using the **tag-protocol-id** statement at the **[edit interfaces interface-name gigether-options ethernet-switch-profile]** hierarchy level.
  - The **tag-protocol-id** statement can be configured only on the first port (port 0) of the PIC. If any other (nonzero) port has the **tag-protocol-id** configuration, the Routing Engine registers an error in the system log and the configuration is ignored.
  - The **tag-protocol-id** statement configured on port 0 of the PIC also applies to the rest of the ports on that PIC.
- The interface *Link Down* event always generates an interrupt; however, the interface *Link Up* event does not generate an interrupt. Therefore, the interface link-up event is detected during the 1-second PIC periodic polling process.
- Generic forward error correction (GFEC) (G.709) and no-FEC modes of operation.
- Diagnostics tools:
  - Line loopback
  - Local loopback
- Fast reroute (FRR)—Based on configurable pre-FEC, bit error rate (BER) is supported and is configured using the **ber-threshold-signal-degrade** statement at the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level.
- *jnx-ifotn.mib* and *otn-mib* as defined in RFC 3591. Note that according to Junos OS security standard, configurable parameters are not supported through SNMP. Only the *get* operation is available through SNMP.
- FEC statistics—corrected errors and corrected error ratio.
- OTN payload pseudorandom binary sequence (PRBS) generation and checking by enabling or disabling PRBS with the **prbs** or **no-prbs** statement at the **[edit interfaces *interface-name* otn-options]** hierarchy level.
- Optical channel data unit (ODU)-level delay measurement.
- At the physical interface level, **flexible-ethernet-service**, **ethernet-ccc**, and **ethernet-tcc** encapsulations are supported. For the **flexible-ethernet-service** encapsulation, the logical level supports **enet2**, **vlan-ccc**, and **vlan-tcc** encapsulations.
- At the logical interface level, **dix**, **vlan-ccc**, and **vlan-tc**c encapsulations are supported.
- Interoperability between 100-Gigabit Ethernet interfaces with CFP transceiver and 100-Gigabit Ethernet interfaces with CFP2 transceiver in LAN PHY framing mode and in OTU4 mode.

The following features are not supported on the P2-100GE-OTN PIC:

- Source MAC learning for accounting
- MAC policing

- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per-queue flow control
- Generic framing procedure-framed (GFP-F) mapping modes over OTN
- General communication channel (GCC)
- OTN interface-level Automatic Protection Switching (APS)
- Insertion, monitoring, and display of OTN header overhead byte
- Black link MIB for integration with transponders
- Optical harness support
- Transport interface and state model (GR-1093)
- Trace tone support
- 15-minute and 1-day performance monitoring counters and historic counters

Layer 2 and Layer 3 Features

The following Layer 2 and Layer 3 features are supported on the P2-100GE-OTN PIC:

- MAC detect link up and link down based on local fault signal or remote fault signal.
- MAC statistics.
- Flow control.
- MAC oversized packet counters based on default MTU value or user-configured MTU value.
- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Per-logical interface source address MAC accounting.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Aggregated Ethernet supports 64 child links that can be configured using the **set chassis aggregated-devices maximum-links** configuration command.
- Maximum of 1024 logical interfaces on an aggregated Ethernet physical interface.
- Support for VLAN tagging, flexible VLAN tagging, and stacked VLAN tagging.
- LACP.
- Link protection.

- 802.3 ah OAM.
- 802.1 ag OAM.
- MPLS FRR.
- SNMP.
- Supports per-VLAN queuing (using Packet Forwarding Engine).

OTN Alarms and Defects

The following OTN alarms and defects are supported on the P2-100GE-OTN PIC:

- LOS—Loss Of Signal
- LOF—Loss Of Frame
- LOM—Loss Of Multiframe
- OTU—Degrade
- OTU—AIS
- OTU—IAE
- OTU—BDI
- OTU—TTIM
- OTU—Signal Degrade
- OTU—Signal Fail
- ODU—Signal Fail
- OTU-FEC—Degrade
- OTU-FEC—Excessive errors
- ODU—Signal Degrade
- ODU—AIS
- ODU—BDI
- ODU—OCI
- ODU—LCK
- ODU—TTIM
- OPU—PTM

TCA Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UAS)

Release History Table

Release	Description
15.1	Starting from Junos OS Release 15.1, you can configure the interfaces on the P2-100GE-OTN PIC on PTX5000 routers, to be a part of the mixed rates and mixed mode aggregated Ethernet bundles.
15.1	Starting from Junos OS Release 15.1, you can configure port-based pseudowire class of service (CoS) classification which includes Layer 3 IPv4, IPv6, and MPLS classification for interfaces with ethernet-ccc encapsulation.

Related Documentation

- Configuring OTN Interfaces on P2-100GE-OTN PIC on page 504

Configuring OTN Interfaces on P2-100GE-OTN PIC

To configure an OTN interface on the P2-100GE-OTN PIC you must configure interface-specific options and OTN-related options for the interface.

To configure the interface-specific options:

1. Go to the **[edit interface *interface-name*]** hierarchy level, where *interface-name* is in the et-*fpc/pic/port* format.

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Configure VLAN tagging on the OTN interface to enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

   ```
   [edit interfaces interface-name ]
   user@host# set vlan-tagging
   ```

3. Configure the maximum transmission unit (MTU) size in bytes for the interface.

   ```
   [edit interfaces interface-name ]
   user@host# set mtu bytes
   ```


4. Configure a VLAN ID for the interface.

   [edit interfaces *interface-name*]
   user@host# **set vlan-id *number***

5. Configure the family for the interface.

   [edit interfaces *interface-name*]
   user@host# **set family *family-name***

6. Configure an IP address for the interface.

   [edit interfaces *interface-name*]
   user@host# **set address *address***

To configure the OTN-related options on the interface:

1. Go to the **[edit interface *interface-name* otn-options]** hierarchy level:

   [edit]
   user@host# **edit interfaces *interface-name* otn-options**

2. Enable the laser on the OTN interface. The laser is disabled by default for all OTN interfaces.

   [edit interfaces *interface-name* otn-options]
   user@host# **set laser-enable**

3. Set an trail trace identifier for the source access point and for the destination access point for ODU and OTU on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi | otu-dapi | otu-expected-receive-dapi | otu-expected-receive-sapi | otu-sapi) *tti-identifier***

4. Ignore the trigger for the defect or set the hold time.

   Configure the hold time for the defect trigger as:

   - *up* with a value—Wait for the hold time delay before clearing the alarm when the defect is absent on the OTN interface.
   - *down* with a value—Wait for the hold time delay before raising the alarm when the defect occurs for the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set trigger (oc-lof | oc-lom | oc-los | oc-tsf | odu-ais | odu-bdi | odu-bei | odu-iae | odu-lck | odu-oci | odu-sd | odu-ttim |opu-ptim | otu-ais | otu-bdi | otu-fec-deg | otu-fec-exe | otu-iae | otu-sd | otu-ttim) (hold-time (down *value* | up *value*) | ignore)**

5. Enable the threshold crossing alarms for the OTN interface along with the trigger for the defect.

   - In Junos OS Release 14.1R2 only:

     [edit interfaces *interface-name* otn-options trigger]
     user@host# **set tca (odu-tca-bbe | odu-tca-es | odu-tca-ses | odu-tca-uas | otu-tca-bbe | otu-tca-es | otu-tca-ses | otu-tca-uas) (enable-tca | no-enable-tca | threshold)**

   - In Junos OS Release 14.2 and later:

     [edit interfaces *interface-name* otn-options]
     user@host# **set tca (odu-tca-bbe | odu-tca-es | odu-tca-ses | odu-tca-uas | otu-tca-bbe | otu-tca-es | otu-tca-ses | otu-tca-uas) (enable-tca | no-enable-tca | threshold)**

6. Set the OTN header bytes as a transmit payload type from 0 bytes through 255 bytes for the packets that are transmitted on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set bytes transmit-payload-type *value***

7. Configure the forward error correction (FEC) mode as Generic Forward Error Correction (GFEC) or none for the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set fec (gfec | none)**

8. Enable line loopback or local host loopback for the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# **set line-loopback**
   user@host# **set local-loopback**

9. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

   [edit interfaces *interface-name* otn-options]
   user@host# **set insert-odu-lck**

10. Enable an ODU open connection indication signal on the OTN interface to send to send the signal pattern 01100110.

    [edit interfaces *interface-name* otn-options]
    user@host# **set insert-odu-oci**

11. Enable a consequent action as listed in the ITU-T G.798 standard for ODU trail trace identifier mismatch (TTIM) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# **set odu-ttim-action-enable**

12. Enable a consequent action as listed in the ITU-T G.798 standard for OTU trail trace identifier mismatch (TTIM) on the OTN interface.

[edit interfaces *interface-name* otn-options]
user@host# **set otu-ttim-action-enable**

13. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# **set prbs**

14. Configure OTN mode as OTU4 for the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# **set rate otu4**

15. Configure the threshold value for signal degradation when an alarm needs to be raised. Configure the threshold value after signal degradation when the alarm needs to be cleared. When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear *value*** statement, then BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

    [edit interfaces *interface-name* otn-options signal-degrade]
    user@host# **set ber-threshold-signal-degrade *value***
    user@host# **set ber-threshold-clear *value***
    user@host# **set interval *value***

16. Enable the following actions for the **preemptive-fast-reroute** statement:

    - Backward FRR—Insert the local pre-FEC status into the transmitted OTN frames and monitor the received OTN frames for the pre-FEC status.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# **set backward-frr-enable**

    - ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# **set odu-backward-frr-enable**

    - Monitoring of signal degradation of pre-FEC OTN frames.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# **set signal-degrade-monitor-enable**

    - Monitoring of signal degradation of ODU BER in the received OTN frames.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# **set odu-signal-degrade-monitor-enable**

17. Configure the following options for ODU BER signal degradation on the OTN interface:

    - Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

[edit interfaces *interface-name* otn-options odu-signal-degrade]
user@host# **set ber-threshold-signal-degrade** ***value***

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set ber-threshold-clear** ***value***

- When you configure the interval along with the **ber-threshold-signal-degrade** ***value*** statement, the ODU bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear** ***value*** statement, then ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set interval** ***value***

Related Documentation

- optics-options on page 1288
- otn-options on page 1289
- signal-degrade on page 971
- preemptive-fast-reroute on page 968

Understanding the MIC3-100G-DWDM MIC

Starting from Junos OS Release 15.1F5, the 100-Gigabit DWDM OTN MIC—MIC3-100G-DWDM—is supported on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG) on the MX240, MX480, MX960, MX2010, and MX2020 routers. The MIC3-100G-DWDM MIC provides a single 100-Gigabit Ethernet interface port that supports DP-QPSK with coherent reception and OTU4 and OTU4 (v) framing modes.

The interfaces on MIC3-100G-DWDM MIC are named with prefix *et*. For more information, see *Interface Naming Overview*.

The following sections explain the features of this MIC in detail:

Interface Features

The following interface features are supported on the MIC3-100G-DWDM MIC:

- Single port 100-Gigabit Ethernet interface with OTU4 (v) framing. DP-QPSK modulation with coherent reception using a CFP2-ACO DWDM optical transceiver.
- Gigabit Ethernet local loopback.
- Diagnostics tools:

- Line loopback
- Local loopback
- Optical Channel Data Unit (ODU) Open Connection Error
- Optical Channel Data Unit (ODU) Lock Maintenance Signal

- Types of forward error corrections (FEC):
  - GFEC (generic forward error correction)
  - HGFEC (high gain forward error correction)
  - SDFEC (soft-decision forward error correction)
- The following MIB modules continue to be supported (and no new MIB is introduced):
  - MIB module to describe Black Link extension to RFC 3591 (jnxoptIfExtMibModule)
  - MIB module to manage the OTN interface (jnxIfOtnMib)
  - MIB module to manage the Optics interface (jnxIfOpticsMib)
  - MIB module to manage OTN FRUs (jnxFruMib)
- Interoperability with the 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM) is not supported.
- Support for interoperability with other vendors' 100 Gigabit Ethernet interfaces.
- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per-queue flow control
- 15-minute and 1-day performance monitoring and historic statistics.
  - Near-end and far-end performance monitoring
  - Threshold-crossing alarms
  - BER performance monitoring
  - FEC performance monitoring
  - Optical performance monitoring
- Insertion, monitoring, and display of OTN header overhead
- Transport interface and state model (GR-1093)

Layer 2 and Layer 3 Features

The following Layer 2 and Layer 3 features are supported on the MIC3-100G-DWDM MIC:

- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Flexible VLAN tagging.
- 802.3 ah OAM.
- 802.1 ag OAM.

OTN Alarms and Defects

The following OTN alarms and defects are supported on the MIC3-100G-DWDM MIC:

Optical Channel(OC) Alarms and Defects

- OC-LOS—Loss Of Signal
- OC-LOF—Loss Of Frame
- OC-LOM—Loss Of Multiframe
- OC-Wavelength-Lock—Wavelength Lock

Optical Channel Data Unit (ODU) Defects

- ODU-AIS—ODU Alarm Indication Signal
- ODU-BDI—ODU Backward Defect Indication
- ODU-BIAE—ODU Backward Incoming Alignment Error
- ODU-IAE—ODU Incoming Alignment Error
- ODU-LCK—ODU Locked
- ODU-LTC—ODU Loss of Tandem Connection
- ODU-OCI—ODU Open Connection Error
- ODU-SSF—ODU Server Signal Failure
- ODU-TSF—ODU Trail Signal Failure
- ODU-TTIM—ODU Trail Trace Identifier Mismatch

Optical Channel Transport Unit (OTU) Defects

- OTU-AIS—OTU Alarm Indication Signal
- OTU-BDI—OTU Backward Defect Indication
- OTU-BIAE—OTU Backward Incoming Alignment Error
- OTU-FEC-DEG—OTU Forward Error Correction Degrade

- OTU-FEC-EXCESS-FEC—OTU Forward Error Correction Excessive FEC Errors
- OTU-IAE—OTU Incoming Alignment Error
- OTU-SSF—OTU Server Signal Failure
- OTU-TSF—OTU Trail Signal Failure
- OTU-TTIM—OTU Trail Trace Identifier Mismatch

Threshold-Crossing Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold —near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minutes interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UES)

Release History Table

Release	Description
15.1F5	Starting from Junos OS Release 15.1F5, the 100-Gigabit DWDM OTN MIC—MIC3-100G-DWDM—is supported on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG) on the MX240, MX480, MX960, MX2010, and MX2020 routers.

Related Documentation

- *Before You Begin Installing or Upgrading the Firmware*
- Configuring OTN Interfaces on MIC3-100G-DWDM MIC on page 511
- Configuring Packet Optical Networks with PTX Series Devices

Configuring OTN Interfaces on MIC3-100G-DWDM MIC

Starting from Junos OS Release 15.1F5, the 100-Gigabit DWDM OTN MIC—MIC3-100G-DWDM—is supported on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG) on the MX240, MX480, MX960, MX2010, and MX2020 routers. To configure an OTN interface on the MIC3-100G-DWDM MIC, you must configure interface-specific options and OTN-related options for the interface.

To configure the interface-specific options:

1. Configure VLAN tagging at the [**edit interface interface-name**] hierarchy level, where interface-name is in the **et-*fpc/pic/port***format.

   [edit interfaces *interface-name*]
   user@host# set vlan-tagging

2. Configure the maximum transmission unit (MTU) size in bytes for the interface.

   [edit interfaces *interface-name*]
   user@host# set mtu *value*

3. Configure a VLAN ID for the interface.

   [edit interfaces *interface-name*]
   user@host# set vlan-id *number*

4. Configure the family for the interface.

   [edit interfaces *interface-name*]
   user@host# set family *family-name*

5. Configure an IP address for the interface.

   [edit interfaces *interface-name*]
   user@host# set address *address*

To configure the optics-specific options on the interface:

1. Specify the optical transmit laser output power in dBm at the [**edit interface** ***interface-name*** **optics-options**] hierarchy level. The default transmit laser output value is 0 dBm.

   [edit interfaces *interface-name* optics-options]
   user@host# set tx-power *value*

2. Specify the wavelength of the optics in nanometers. For a list of wavelengths supported, see wavelength.

   [edit interfaces *interface-name* optics-options]
   user@host# set wavelength nm

To configure the OTN-specific options on the interface:

1. At the [**edit interfaces** ***interface-name*** **otn-options**] enable the laser on the OTN interface. The laser is disabled by default for all OTN interfaces.

   [edit interfaces *interface-name* otn-options]
   user@host# set laser-enable

2. Set an trail trace identifier for the source access point and for the destination access point for ODU and OTU on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# set tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi | otu-dapi | otu-expected-receive-dapi | otu-expected-receive-sapi | otu-sapi)

3. By default, triggers are ignored. Specify defect triggers and the set the trigger hold time for the trigger. Possible values for the trigger hold time are as follows: down—Delay

before marking interface down when defect occurs (1..65534 milliseconds) and up—Delay before marking interface up when defect is absent (1..65534 milliseconds).

NOTE: The hold time value only impacts the alarm reporting time and does not mark an interface down when the defect occurs. To mark the interface up or down, you must also configure the physical interface hold time at the [edit interfaces *interface-name*] hierarchy level.

[edit interfaces *interface-name* otn-options]
user@host# set trigger (oc-lof | oc-lom | oc-los | oc-tsf | odu-ais | odu-bdi | odu-bei | odu-iae | odu-lck | odu-oci | odu-sd | odu-ttim |opu-ptim | otu-ais | otu-bdi | otu-fec-deg | otu-fec-exe | otu-iae | otu-sd | otu-ttim) (hold-time (down value | up value) | ignore)

4. Enable the threshold crossing alarms for the OTN interface along with the trigger for the defect.

   [edit interfaces *interface-name* otn-options]
   user@host# set tca (odu-tca-bbe | odu-tca-es | odu-tca-ses | odu-tca-uas | otu-tca-bbe | otu-tca-es | otu-tca-ses | otu-tca-uas) (enable-tca | no-enable-tca | threshold)

5. Set the OTN header bytes as a transmit payload type from 0 bytes through 255 bytes for the packets that are transmitted on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# set bytes transmit-payload-type value

6. Configure the forward error correction (FEC) mode for the OTN interface. Possible values are: Generic Forward Error Correction (GFEC), or High Gain Forward Error Correction (HGFEC) or Soft Decision Forward Error Correction (SDFEC). The default forward error correction mode is SDFEC.

   [edit interfaces *interface-name* otn-options]
   user@host# set fec (gfec | hgfec | sdfec)

7. Enable line loopback or local host loopback for the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# set line-loopback
   user@host# set local-loopback

8. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

   [edit interfaces *interface-name* otn-options]
   user@host# set insert-odu-lck

9. Enable an ODU open connection indication signal on the OTN interface to send to send the signal pattern 01100110.

   [edit interfaces *interface-name* otn-options] user@host# set insert-odu-oci

10. Enable a consequent action as listed in the ITU-T G.798 standard for ODU trail trace identifier mismatch (TTIM) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# set odu-ttim-action-enable

11. Enable a consequent action as listed in the ITU-T G.798 standard for OTU trail trace identifier mismatch (TTIM) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# set out-ttim-action-enable

12. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# set prbs

13. Configure the line rate or speed of the OTN signal to OTU4 (100Gbps) for the OTN interface.

    NOTE: If you specify a value other than OTU4, the value is ignored. To verify the line rate, use the **show interfaces *interface-name* extensive** command.

    [edit interfaces *interface-name* otn-options]
    user@host# set rate otu4

14. Configure the threshold value for signal degradation when an alarm needs to be raised. Configure the threshold value after signal degradation when the alarm needs to be cleared. When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear *value*** statement, then BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

    [edit interfaces *interface-name* otn-options signal-degrade]
    user@host# set ber-threshold-signal-degrade value
    user@host# set ber-threshold-clear value
    user@host# set interval value

15. Enable the following actions for the preemptive-fast-reroute statement:

    - Backward FRR—Insert the local pre-FEC status into the transmitted OTN frames and monitor the received OTN frames for the pre-FEC status.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# set backward-frr-enable

- ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

  [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
  user@host# **set odu-backward-frr-enable**

- Monitoring of signal degradation of pre-FEC OTN frames.

  [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
  user@host# **set signal-degrade-monitor-enable**

- Monitoring of signal degradation of ODU BER in the received OTN frames.

  [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
  user@host# **set odu-signal-degrade-monitor-enable**

16. Configure the following options for ODU BER signal degradation on the OTN interface:

- Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set ber-threshold-signal-degrade *value***

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set ber-threshold-clear *value***

- When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the ODU bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear *value*** statement, then ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set interval *value***

Related Documentation

- Understanding the MIC3-100G-DWDM MIC on page 508
- **optics-options on page 1288**
- **otn-options on page 1289**
- **signal-degrade on page 971**
- **preemptive-fast-reroute on page 968**

Understanding the PTX-5-100G-WDM PIC

Starting with Junos OS Release 15.1F6, the 5-port 100-Gigabit DWDM OTN PIC—PTX-5-100G-WDM—is supported on the PTX3000 and the PTX5000 routers. The PTX-5-100G-WDM PIC provides five 100-Gigabit Ethernet interface ports that support dual-polarization quadrature phase shift keying (DP-QPSK) modulation with coherent reception and OTU4 and OTU4 (v) framing modes.

The interfaces on the PTX-5-100G-WDM PIC are named with the prefix *et*. For more information, see *Interface Naming Overview*.

NOTE: **The 5-port 100-Gigabit DWDM OTN PIC is not directly interoperable with the 2-port 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM), but they can both operate over the same DWDM line system.**

The following sections explain the features of this PIC in detail:

Interface Features

The following interface features are supported on the PTX-5-100G-WDM PIC:

- Five-port 100-Gigabit Ethernet interface with OTU4 (v) framing and DP-QPSK modulation with coherent reception using a CFP2-ACO DWDM optical transceiver.
- Gigabit Ethernet local loopback.
- Diagnostics tools:
  - Line loopback
  - Local loopback
  - Optical Channel Data Unit (ODU) Open Connection Error
  - Optical Channel Data Unit (ODU) Lock Maintenance Signal
- Types of forward error corrections (FEC):
  - GFEC (generic forward error correction)

    NOTE: **GFEC mode is not supported on Junos OS Release 15.1F6. Junos OS Release 15.1F6-S1 supports GFEC mode. Contact customer support for the Junos OS Release 15.1F6-S1.**

  - SDFEC (soft-decision forward error correction)
- The following MIB features continue to be supported (and no new MIB is introduced):

  - MIB module to describe Black Link extension to RFC 3591 (jnxoptIfExtMibModule). The Black Link extension enables an optical transceiver of a vendor to introduce an optical signal over an optical network from another vendor.
  - MIB module to manage the OTN interface (jnxIfOtnMib)
  - MIB module to manage the Optics interface (jnxIfOpticsMib)
  - MIB module to manage OTN FRUs (jnxFruMib)
- Interoperability with other vendors' 100 Gigabit-Ethernet interfaces.
- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per-queue flow control
- 15-minute and 1-day performance monitoring and historic statistics.
  - Near-end and far-end performance monitoring
  - Threshold-crossing alarms
  - BER performance monitoring
  - FEC performance monitoring
  - Optical performance monitoring
- Insertion, monitoring, and display of OTN header overhead
- Transport interface and state model (GR-1093)

Layer 2 and Layer 3 Features

The following Layer 2 and Layer 3 features are supported on the PTX-5-100G-WDM PIC:

- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Flexible VLAN tagging.
- 802.3 ah OAM.
- 802.1 ag OAM.

OTN Alarms and Defects

The following OTN alarms and defects are supported on the PTX-5-100G-WDM PIC:

Optical Channel Alarms and Defects

- OC-LOS—Loss Of Signal
- OC-LOF—Loss Of Frame
- OC-LOM—Loss Of Multiframe
- OC-Wavelength-Lock—Wavelength Lock

Optical Channel Data Unit (ODU) Defects

- ODU-AIS—ODU Alarm Indication Signal
- ODU-BDI—ODU Backward Defect Indication
- ODU-BIAE—ODU Backward Incoming Alignment Error
- ODU-IAE—ODU Incoming Alignment Error
- ODU-LCK—ODU Locked
- ODU-LTC—ODU Loss of Tandem Connection
- ODU-OCI—ODU Open Connection Error
- ODU-SSF—ODU Server Signal Failure
- ODU-TSF—ODU Trail Signal Failure
- ODU-TTIM—ODU Trail Trace Identifier Mismatch

Optical Channel Transport Unit (OTU) Defects

- OTU-AIS—OTU Alarm Indication Signal
- OTU-BDI—OTU Backward Defect Indication
- OTU-BIAE—OTU Backward Incoming Alignment Error
- OTU-FEC-DEG—OTU Forward Error Correction Degrade
- OTU-FEC-EXCESS-FEC—OTU Forward Error Correction Excessive FEC Errors
- OTU-IAE—OTU Incoming Alignment Error
- OTU-SSF—OTU Server Signal Failure
- OTU-TSF—OTU Trail Signal Failure
- OTU-TTIM—OTU Trail Trace Identifier Mismatch

Threshold Crossing Alarms

Threshold-crossing alarms (TCAs) are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15-minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UES)

Related Documentation

- *Before You Begin Installing or Upgrading the Firmware*
- Configuring OTN Interfaces on PTX-5-100G-WDM PIC on page 519
- *Installing Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*
- *Upgrading Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*
- Configuring Packet Optical Networks with PTX Series Devices

Configuring OTN Interfaces on PTX-5-100G-WDM PIC

Starting from Junos OS Release 15.1F6, the the 5-port 100-Gigabit DWDM OTN PIC—PTX-5-100G-WDM—is supported on the PTX3000 and the PTX5000 routers. To configure an OTN interface on the PTX-5-100G-WDM PIC, you must configure interface-specific options, optics-specific options and OTN-related options for the interface.

To configure the interface-specific options:

1. Configure VLAN tagging at the [**edit interface *interface-name***] hierarchy level, where *interface-name* is in the **et-*fpc/pic/port*** format.

   ```
   [edit interfaces interface-name]
   user@host# set vlan-tagging
   ```

2. Configure the maximum transmission unit (MTU) size in bytes for the interface. Possible values: 256 through 16,000.

   ```
   [edit interfaces interface-name]
   user@host# set mtu value
   ```

3. Set the unit VLAN ID, family and the IP address of the interface. Possible values for the VLAN ID: 1 through 4094. Specify the family as **inet**.

   ```
   [edit interfaces interface-name unit 0]
   user@host# set vlan-id number
   user@host# set family family-name
   user@host# set address address
   ```

To configure the optics-specific options on the interface:

1. Specify the optical transmit laser output power in dBm at the [**edit interface *interface-name* optics-options**] hierarchy level. The default transmit laser output value is 0 dBm.

   [edit interfaces *interface-name* optics-options]
   user@host# set tx-power *value*

2. Specify the wavelength of the optics in nanometers. For a list of wavelengths supported, see **wavelength**.

   [edit interfaces *interface-name* optics-options]
   user@host# set wavelength nm

To configure the OTN-specific options on the interface:

1. At the [**edit interfaces *interface-name* otn-options**] hierarchy level, enable the laser on the OTN interface. The laser is disabled by default for all OTN interfaces.

   [edit interfaces *interface-name* otn-options]
   user@host# set laser-enable

2. Set a trail trace identifier for the source access point and for the destination access point for ODU and OTU on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# set tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi | otu-dapi | otu-expected-receive-dapi | otu-expected-receive-sapi | otu-sapi)

3. Specify defect triggers and the set the trigger hold time for the trigger. By default, triggers are ignored. Possible values for the trigger hold time are as follows: down and up.

   - down—Delay before marking interface down when defect occurs (1 through 65534 milliseconds)
   - up—Delay before marking interface up when defect is absent (1 through 65534 milliseconds).

   NOTE: The hold time value only impacts the alarm reporting time and does not mark an interface down when the defect occurs. To mark the interface up or down, you must also configure the physical interface hold time at the [**edit interfaces *interface-name***] hierarchy level.

   [edit interfaces *interface-name* otn-options]
   user@host# set trigger (oc-lof | oc-lom | oc-los | oc-tsf | odu-ais | odu-bdi | odu-bei | odu-iae | odu-lck | odu-oci | odu-sd | odu-ttim |opu-ptim | otu-ais | otu-bdi | otu-fec-deg | otu-fec-exe | otu-iae | otu-sd | otu-ttim) (hold-time (down *value* | up *value*) | ignore)

4. Enable the threshold-crossing alarms (TCAs) for the OTN interface along with the trigger for the defect. Threshold-crossing alarms (TCAs) are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15-minute interval for parameters such as OTU and ODU.

   [edit interfaces *interface-name* otn-options]
   user@host# set tca (odu-tca-bbe | odu-tca-es | odu-tca-ses | odu-tca-uas | otu-tca-bbe | otu-tca-es | otu-tca-ses | otu-tca-uas) (enable-tca | no-enable-tca | threshold)

5. Set the OTN header bytes as a transmit payload type from 0 bytes through 255 bytes for the packets that are transmitted on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# set bytes transmit-payload-type value

6. Configure the forward error correction (FEC) mode for the OTN interface. Possible values are: generic forward error correction (GFEC), or high-gain forward error correction (HG-FEC) or soft-decision forward error correction (SD-FEC). The default forward error correction mode is SD-FEC.

   [edit interfaces *interface-name* otn-options]
   user@host# set fec (gfec |hgfec | sdfec)

7. Enable line loopback or local host loopback for the OTN interface. Loopback testing enables you to verify the connectivity of a circuit. In line loopback, instead of transmitting the signal toward the far-end device, the signal is sent back to the originating router. In local loopback, the signal is transmitted to the channel service unit (CSU) and then to the far-end device.

   [edit interfaces *interface-name* otn-options]
   user@host# set line-loopback
   user@host# set local-loopback

8. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

   [edit interfaces *interface-name* otn-options]
   user@host# set insert-odu-lck

9. Enable an ODU open connection indication signal on the OTN interface to send the signal pattern 01100110.

   [edit interfaces *interface-name* otn-options]
   user@host# set insert-odu-oci

10. Enable a consequent action as listed in the ITU-T G.798 standard for ODU trail trace identifier mismatch (TTIM) on the OTN interface.

   [edit interfaces *interface-name* otn-options]
   user@host# set odu-ttim-action-enable

11. Enable a consequent action as listed in the ITU-T G.798 standard for OTU trail trace identifier mismatch (TTIM) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# set out-ttim-action-enable

12. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

    [edit interfaces *interface-name* otn-options]
    user@host# set prbs

13. Configure the line rate or speed of the OTN signal to otu4 (100 Gbps) for the OTN interface.

    NOTE: If you specify a value other than otu4, the value is ignored. To verify the line rate, use the **show interfaces *interface-name* extensive** command.

    [edit interfaces *interface-name* otn-options]
    user@host# set rate otu4

14. Configure the threshold value for signal degradation when an alarm needs to be raised. Configure the threshold value after signal degradation when the alarm needs to be cleared. When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear *value*** statement, then BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

    [edit interfaces *interface-name* otn-options signal-degrade]
    user@host# set ber-threshold-signal-degrade *value*
    user@host# set ber-threshold-clear *value*
    user@host# set interval *value*

15. Enable the following actions for the **preemptive-fast-reroute** statement:

    - Backward FRR—Insert the local pre-FEC status into the transmitted OTN frames and monitor the received OTN frames for the pre-FEC status.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# set backward-frr-enable

    - ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# **set odu-backward-frr-enable**

    - Monitoring of signal degradation of pre-FEC OTN frames.

      [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
      user@host# **set signal-degrade-monitor-enable**

- Monitoring of signal degradation of ODU BER in the received OTN frames.

  [edit interfaces *interface-name* otn-options preemptive-fast-reroute]
  user@host# **set odu-signal-degrade-monitor-enable**

16. Configure the following options for ODU BER signal degradation on the OTN interface:

- Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set ber-threshold-signal-degrade** ***value***

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set ber-threshold-clear** ***value***

- When you configure the interval along with the **ber-threshold-signal-degrade** ***value*** statement, the ODU bit error rate (BER) must stay above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the **ber-threshold-clear** ***value*** statement, then ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

  [edit interfaces *interface-name* otn-options odu-signal-degrade]
  user@host# **set interval** ***value***

Release History Table

Release	Description
15.1F6	Starting from Junos OS Release 15.1F6, the the 5-port 100-Gigabit DWDM OTN PIC—PTX-5-100G-WDM—is supported on the PTX3000 and the PTX5000 routers.

Related Documentation

- *Before You Begin Installing or Upgrading the Firmware*
- *Installing Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*
- Understanding the PTX-5-100G-WDM PIC on page 516
- *Upgrading Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*
- **optics-options on page 1288**
- **otn-options on page 1289**
- **signal-degrade on page 971**
- **preemptive-fast-reroute on page 968**

Understanding ODU Path Delay Measurement on OTN Networks for Performance Monitoring

Performance monitoring is an important requirement in any network, including the optical transport networks (OTN). The key parameters that impact performance are bit error rate (BER) and delay. Delays in data communication over a network impact the network latency. Network latency is the time taken for a packet of data to travel from a designated point to another designated point. If there are less delays, the network latency is low. You can measure latency by sending a packet and then receiving it as it is returned back to you; the time taken for the round-trip indicates the latency.

The optical channel data unit (ODU) path delay measurement offers in-service delay measurement. Delay (or latency) is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (**DM**) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end. For instance, if the transmitted delay measurement bit is **1 1 1 1 1 1 1 1 0 0** and the received delay measurement bit is **1 1 1 0 0 0 0 0 0 0**, the delay measurement starts at frame 2 and ends at frame 8. This can be detected by the change in value between the transmitted bit and the received bit.

```
Frame#   10 9 8 7 6 5 4 3 2 1
Tx DM bit  1 1 1 1 1 1 1 1 0 0
Rx DM bit  1 1 1 0 0 0 0 0 0 0
```

The result of the delay measurement is 6 frames (8 - 2).

Guidelines for Configuring Delay Measurement

When you configure in-service delay measurement, we recommend that you follow certain guidelines to ensure that you obtain accurate delay measurement.

- Unidirectional delay measurement is not supported. The in-service delay measurement is specific to round-trip delay measurement and for optical channel data units only.
- Delay measurement on different framers for the MIC and PIC is different. So, the delay measurement values are different.
- Resiliency is not supported for path delay measurement.
- Links at the local and remote interfaces must be active before you configure delay measurement.
- Do not perform delay measurement tests when ODU maintenance signals are injected.
- Do not configure local loopback and network loopback with remote loopback because the loopback data is overwritten by the delay measurement pattern.

NOTE: If a link failure occurs after you begin measuring delay, delay measurement fails. You must re-enable measurement of delay on the local interface to measure delay.

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- *100-Gigabit DWDM OTN MIC with CFP2-ACO*
- *100-Gigabit DWDM OTN PIC with CFP2-ACO (PTX Series)*
- Configuring OTN Interfaces on MIC3-100G-DWDM MIC on page 511
- Configuring OTN Interfaces on PTX-5-100G-WDM PIC on page 519
- Disabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 527
- Enabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 525
- **remote-loop-enable on page 970**
- Understanding the MIC3-100G-DWDM MIC on page 508
- Understanding the PTX-5-100G-WDM PIC on page 516

Enabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring

Delay measurement is disabled by default. This topic explains the broad steps for measuring the optical channel data units (ODU) path delay on optical transport networks (OTN). First, enable remote loopback on the remote interface and commit the configuration. This enables the remote interface to loop back the delay measurement pattern to the local interface. Then, start delay measurement at the local interface and view the results.

NOTE: **Do not enable remote loopback on both ends (local and remote). If you enable remote loopback on both interfaces, the delay measurement pattern is looped back continuously between the two interfaces.**

Before you start measuring delay in the ODU path on OTN, complete the following tasks:

- Ensure that the links are active at the local and remote interfaces and alarms are not configured.
- Ensure that there is a delay of 10 seconds before enabling remote loopback. Also, ensure that there is a delay of 10 seconds after enabling remote loopback at the remote interface and before you start measuring delay.
- Ensure that the delay measurement tests are not performed when ODU maintenance signals are injected.
- Ensure that the local loopback and network loopback are also not specified because the looped-back data is overwritten by the delay measurement pattern.

NOTE: If link failure occurs after you begin measuring delay, delay measurement fails. You must re-enable measurement of delay on the local interface to measure delay.

To enable ODU path delay measurement, first enable remote loopback of the delay measurement pattern on the remote interface and then start measurement of the delay.

1. Enable remote loopback on the remote interface by including the **remote-loop-enable** statement at the [**edit**] hierarchy level.

   [edit]
   user@host# set interfaces *interfacename* otn-options odu-delay-management remote-loop-enable

2. After enabling remote loopback, commit the configuration.

   [edit]
   user@host# commit

3. Start delay measurement on the local interface by including the **start-measurement** statement at the [**edit**] hierarchy level.

   [edit]
   user@host# set interface *interfacename* otn-options odu-delay-management start-measurement

4. After enabling measurement of delay on the local interface, commit the configuration.

   [edit]
   user@host# commit

5. To view the delay measurement values, from the operational mode, enter the **show interfaces extensive** command.

   user@host> show interfaces *interfacename* extensive

```
...
ODU Delay Management:
Start Measurement: True
Remote Loop Enable: False
Result: 0 micro seconds
...
```

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- *100-Gigabit DWDM OTN MIC with CFP2-ACO*
- *100-Gigabit DWDM OTN PIC with CFP2-ACO (PTX Series)*
- Configuring OTN Interfaces on MIC3-100G-DWDM MIC on page 511
- Configuring OTN Interfaces on PTX-5-100G-WDM PIC on page 519

- Disabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 527
- **remote-loop-enable on page 970**
- Understanding ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 524
- Understanding the MIC3-100G-DWDM MIC on page 508
- Understanding the PTX-5-100G-WDM PIC on page 516

Disabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring

Delay measurement is disabled by default. If you enabled optical channel data unit (ODU) path delay measurement by using the **remote-loop-enable** and **start-measurement** statements, you can use this procedure to disable delay measurement.

NOTE: You can also use the **delete** or **deactivate** command to disable remote loopback on the remote interface. For instance, you can use the **delete interfaces *interfacename* otn-options odu-delay-management remote-loop-enable** or **deactivate interface *interfacename* otn-options odu-delay-management remote-loop-enable** command to disable remote loopback on the remote interface.

To disable ODU path delay measurement, first disable remote loopback of the delay measurement pattern on the remote interface and then stop delay measurement:

1. Stop delay measurement on the local interface by including the **stop-measurement** statement at the [**edit**] hierarchy level.

   ```
   [edit]
   user@host# set interface interfacename otn-options odu-delay-management
     stop-measurement
   ```

2. After you stop delay measurement on the local interface, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

3. Disable remote loopback on the remote interface by including the **no-remote-loop-enable** statement at the [**edit**] hierarchy level.

   ```
   [edit]
   user@host# set interfaces interfacename otn-options odu-delay-management
     no-remote-loop-enable
   ```

4. After disabling remote loopback on the remote interface, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```


5. To verify that remote loopback is disabled and delay is not measured, enter the **show interfaces extensive** command, from the operational mode.

   user@host> show interfaces *interfacename* extensive

```
...
ODU Delay Management:
Start Measurement: False
Remote Loop Enable: False
Result: 0 micro seconds
...
```

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- *100-Gigabit DWDM OTN MIC with CFP2-ACO*
- *100-Gigabit DWDM OTN PIC with CFP2-ACO (PTX Series)*
- Configuring OTN Interfaces on MIC3-100G-DWDM MIC on page 511
- Configuring OTN Interfaces on PTX-5-100G-WDM PIC on page 519
- Enabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 525
- **remote-loop-enable on page 970**
- Understanding ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 524
- Understanding the MIC3-100G-DWDM MIC on page 508
- Understanding the PTX-5-100G-WDM PIC on page 516

CHAPTER 28

Configuring Gigabit Ethernet Accounting and Policing

Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs

For Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure granular per-VLAN class-of-service (CoS) capabilities and extensive instrumentation and diagnostics on a per-VLAN and per-MAC address basis.

VLAN rewrite, tagging, and deleting enables you to use VLAN address space to support more customers and services.

VPLS allows you to provide a point-to-multipoint LAN between a set of sites in a VPN. Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) are combined with VPLS to deliver metro Ethernet service.

For Gigabit Ethernet IQ2 and IQ2-E and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, you can apply Layer 2 policing to logical interfaces in the egress or ingress direction. Layer 2 policers are configured at the **[edit firewall]** hierarchy level. You can also control the rate of traffic sent or received on an interface by configuring a policer overhead at the **[edit chassis fpc *slot-number* pic *slot-number*]** hierarchy level.

Table 50 on page 530 lists the capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Table 50: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
Layer 2		
802.3ad link aggregation	Yes	Yes
Maximum VLANs per port	384	1023
Maximum transmission unit (MTU) size	9192	9192
MAC learning	Yes	Yes
MAC accounting	Yes	Yes
MAC filtering	Yes	Yes
Destinations per port	960	960
Sources per port	64	64
Hierarchical MAC policers	Yes, premium and aggregate	No, aggregate only
Multiple TPID support and IP service for nonstandard TPIDs	Yes	Yes
Multiple Ethernet encapsulations	Yes	Yes
Dual VLAN tags	Yes	No
VLAN rewrite	Yes	No
Layer 2 VPNs		
VLAN CCC	Yes	Yes
Port-based CCC	Yes	Yes
Extended VLAN CCC Virtual Metropolitan Area Network (VMAN) Tag Protocol	Yes	Yes
CoS		
PIC-based egress queues	Yes	Yes
Queued VLANs	Yes	No

Table 50: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (continued)

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
VPLS	Yes	Yes

For more information about configuring VPLS, see the *Junos OS VPNs Library for Routing Devices*.

You can also configure CoS on logical IQ interfaces. For more information, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 546
- Configuring MAC Address Accounting on page 531
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring MAC Address Accounting

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, for 100-Gigabit Ethernet Type 5 PIC with CFP, and for MPC3E, MPC4E, MPC5E, MPC5EQ, and MPC6E MPCs, you can configure whether source and destination MAC addresses are dynamically learned.

To configure MAC address accounting on an individual Ethernet interface, include the **mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
mac-learn-enable;
```

To configure MAC address accounting on an aggregated Ethernet interface, include the **mac-learn-enable** statement at the **[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces aex aggregated-ether-options ethernet-switch-profile]
mac-learn-enable;
```

To prohibit an interface from dynamically learning source and destination MAC addresses, do *not* include the **mac-learn-enable** statement.

To disable dynamic learning of the source and destination MAC addresses after it has been configured, you must delete **mac-learn-enable** from the configuration.

NOTE: MPCs support MAC address accounting for an individual interface or an aggregated Ethernet interface member link only after the interface has received traffic from the MAC source. If traffic is only exiting an interface, the MAC address is not learned and MAC address accounting does not occur.

Related Documentation

- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 529
- Configuring Gigabit Ethernet Policers on page 539
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 546
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces Overview

Junos OS supports the capability to compute MAC address statistics for dynamically learned static and destination MAC addresses on physical interfaces. Starting in Junos OS Release 15.1, Junos OS enables you to configure source MAC (SMAC) address and destination MAC (DMAC) address-based accounting for MAC addresses that are dynamically learned on aggregated Ethernet (ae-) interfaces in routed mode. When you include the **mac-learn-enable** statement at the **[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]** hierarchy level, dynamic learning of source and destination MAC addresses is enabled. By default, this capability is disabled. When dynamic learning of MAC addresses is enabled for AE interfaces in routed mode, the MAC-filter settings are updated for each of the child links of the AE bundle interface. This feature provides for both the configuration of the mac-learn-enable filter and the display of SMAC and DMAC based accounting information on the aggregated interface in the output of the **show interfaces mac-database *interface-name* mac-address *mac-address*** command.

When this functionality is enabled, source and destination MAC addresses-based accounting is supported on the routed interfaces on MX Series routers with DPCs and MPCs. Support for mixed mode LAG interfaces is also available. This feature supports MAC address accounting for AE interfaces in routed mode (for inet family). Destination MAC-based accounting is supported only for MAC addresses dynamically learned at the ingress interface, including each individual child or member link of the AE bundle. This behavior occurs because MPCs do not support destination MAC address learning. As a result, if a packet exits a child link without passing in the ingress direction through that link, destination MAC (DMAC) accounting for this packet occurs at the child link level and this data is not available at the aggregate level. Dynamic learning of MAC addresses can be supported on only the AE interface or on selective individual member links. MAC learning support on the bundle depends on the capability of individual member links. If a link in the bundle does not contain the capability to support MAC learning or accounting, it is disabled on the AE bundle.

The MAC data for the aggregated bundle is displayed by collecting data form individual child links. This data is collected when the command to display the MAC database is triggered from the CLI. This method of data collection implies that based on the number of child links and the size of the MAC database, the time take to display the database differs. This approach to obtain the current snapshot of the MAC database from the currently active child links is used instead of maintaining a database at the Routing Engine because of the dynamic nature of the MAC database and the overhead required to maintain the database information in synchronization with all the child Packet Forwarding Engines. A difference in the DMAC-based accounting for packets generated from the Routing Engine (packets sent in the host path). On DPCs, these packets are accounted in egress direction (Output Packet/Byte count), whereas on MPCs, these packets are not accounted because DMAC learning is not supported. This difference in behavior also occurs between child links on DPCs and MPCs. Because this feature to enable dynamic learning is related to collecting MAC database statistics from child links based on the command issued from the CLI, there is a impact on the time it takes to display the data on the console based on the size of the MAC database and the number of child-links spread across different FPCs. The limit on the maximum number of MAC addresses that can be learned from an interface does not apply to this dynamic learning of MAC addresses functionality.

Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, Junos OS enables you to configure source MAC (SMAC) address and destination MAC (DMAC) address-based accounting for MAC addresses that are dynamically learned on aggregated Ethernet (ae-) interfaces in routed mode.

Related Documentation

- mac-learn-enable on page 1245

Accounting of the Layer 2 Overhead Attribute in Interface Statistics

On MX Series and T Series routers, you can configure the logical interface statistics to include the Layer 2 overhead size (header and trailer bytes) for both ingress and egress interfaces. Both the transit and total statistical information are computed and displayed for each logical interface. This functionality is supported on 1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces on Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs) on MX Series routers. Starting with Junos OS Release 13.2, configuring the logical interface statistics to include Layer 2 is supported on 10-Gigabit Ethernet interfaces on MX Series routers with MPC4E. Starting with Junos OS Release 13.3, **account-layer2-overhead** is not supported on MX Series routers with MPC3E (on both PIC and logical interface levels).

You can also configure the capability to compute the Layer 2 overhead bytes in interface statistics on Type-3, Type-4 and Type-5 Flexible Port Concentrators (FPCs) on T Series routers. To enable the Layer 2 overhead bytes to be counted in the interface statistics at the PIC level, you must use the **account-layer2-overhead** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level.

If you configure this capability, all the Layer 2 header details (Layer 2 header and cyclic redundancy check [CRC]) based on the Layer 2 encapsulation configured for an interface are calculated and displayed in the logical interface statistics for ingress and egress interfaces in the output of the **show interfaces** ***interface-name*** commands. For logical interfaces, the **Input bytes** and **Output bytes** fields under the Traffic statistics section in the output of the **show interfaces** ***interface-name*** **<detail | extensive>** command include the Layer 2 overhead of the packets. For logical interfaces, the Input rate and Output rate fields under the Traffic statistics section in the output of the **show interfaces** ***interface-name*** **<media | statistics>** command include the Layer 2 overhead of the packets. For logical interfaces, the values for the newly added **Egress account overhead** and **Ingress account overhead** fields display the Layer 2 overhead size for transmitted and received packets respectively.

The input and output octets at the logical interface configured on the PIC includes all the Layer 2 headers. All the logical interfaces on the PIC, including the ae and the non-ae interfaces, are processed for Layer 2 overhead accounting for the arriving and exiting packets. This method of operation impacts the transit statistics that are primarily used for subscriber accounting and billing purposes in customer networks.

Table 51 on page 534 lists the adjustment bytes that are counted based on the encapsulation on the logical interface over the Ethernet interface, when you enable accounting of Layer 2 overhead in interface statistics at the PIC level. The values for the adjustment bytes that are listed for all types of encapsulation are the same for DPCs and MPCs, with the only exception being for the VLAN CCC adjustment value. On DPCs, the VLAN CCC adjustment value is –4 bytes and on MPCs, the VLAN CCC adjustment value is +4 bytes.

Table 51: Adjustment Bytes for Logical Interfaces over Ethernet Interfaces

Encapsulation Type on Logical Interfaces	Number of Adjustment Bytes	Description
Ethernet DIXv2 (IP datagrams over Ethernet)	18	Untagged (includes CRC)
Ethernet DIXv2 (IP datagrams over Ethernet)	22	Single-tagged (includes CRC)
Ethernet DIXv2 (IP datagrams over Ethernet)	26	Double-tagged (includes CRC)
VLAN Bridge	4	CRC
VLAN CCC	4	CRC
VLAN TCC	18	Untagged (includes CRC)
VLAN TCC	22	Single-tagged (includes CRC)
VLAN TCC	26	Double-tagged (includes CRC)
VLAN VPLS	4	CRC

Guidelines for Configuring the Computation of Layer 2 Overhead in Interface Statistics

Keep the following points in mind when you configure the computation of Layer 2 overhead in interface statistics:

- When you configure a native VLAN ID on a logical interface, the Layer 2 header adjustment for input statistics is different for tagged and untagged packets. For such interfaces, if you configure the setting to account for Layer 2 overhead, incorrect statistics might be displayed.
- An untagged packet is considered as a tagged packet and an additional 4 bytes are appended to the counter values displayed in the output of the **show interface** command.
- The computed statistics might not be completely accurate in scenarios where the packets are dropped after they have been included in the interface statistics, but before the packets reach the destination.
- Label-switched interface (LSI) statistics on the ingress direction of interfaces do not include the Layer 2 overhead bytes because this functionality of accounting Layer 2 overhead is not supported for such LSI interfaces.
- Layer 2 overhead accounting is not supported for inline service (si) interfaces.
- The total statistics of interfaces do not indicate the complete Layer 2 adjusted statistics. This behavior occurs because the total statistics count is the sum of transit and local statistics. Only the transit statistics are adjusted for Layer 2 and the local statistics are not adjusted for Layer 2.
- Statistics on ae interfaces are calculated in the same manner as non-ae interfaces.
- Adjustment bytes are applicable only for transit statistics that are displayed for logical interfaces.
- For physical interfaces, the adjustment bytes for transit traffic and the non-adjusted bytes for local or protocol-specific traffic are combined and displayed in the output of the **show interfaces** command. (Segregation is not possible.)
- Layer 2 overhead accounting can be enabled at both PIC level and logical interface level.
- When the **account-layer2-overhead** statement is configured, the Layer 2 overhead size in both input and output statistics is accounted for in Dense Port Concentrator (DPCs) and Modular Port Concentrator (MPCs).
- This **account-layer2-overhead** configuration now supports Layer 2 accounting for the Ethernet bridge encapsulation.
- The Layer 2 overhead bytes in interface statistics are saved across a unified ISSU or a graceful Routing Engine switchover (GRES) operation.

Release History Table

Release	Description
13.3	Starting with Junos OS Release 13.3, **account-layer2-overhead** is not supported on MX Series routers with MPC3E (on both PIC and logical interface levels).
13.2	Starting with Junos OS Release 13.2, configuring the logical interface statistics to include Layer 2 is supported on 10-Gigabit Ethernet interfaces on MX Series routers with MPC4E.

Related Documentation

- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 536
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 537
- **account-layer2-overhead on page 1067**

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Layer 2 Overhead Accounting in Interface Statistics

This topic contains sections that describe the configuration of Layer 2 overhead accounting for interface statistics at the PIC level and logical interface level.

Layer 2 overhead accounting can be enabled at both PIC level and logical interface level through configuration. By default, the physical interface and logical interface statistics do not account for Layer 2 overhead size (header and trailer) in both input and output statistics.

When the **account-layer2-overhead** statement is configured, the Layer 2 overhead size in both input and output statistics is accounted for in the Dense Port Concentrator (DPCs) and the Modular Port Concentrator (MPCs). This **account-layer2-overhead** configuration now supports Layer 2 accounting for the Ethernet bridge encapsulation.

Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level

You can configure the **account-layer2-overhead** statement at the **edit chassis fpc *slot-number* pic *pic-number*** hierarchy level to enable accounting of Layer 2 overhead bytes in the ingress and egress interface statistics at the PIC level.

CAUTION: **If you modify the setting for accounting of Layer 2 overhead bytes at the PIC level, the PIC is rebooted, causing all of the physical and logical interfaces to be deleted and readded on the PIC. Due to this behavior, we recommend that you exercise caution while using this feature.**

The computation method of Layer 2 overhead on different interface types is as follows:

- For Ethernet interfaces, all the Layer 2 headers are counted.

- For non-Ethernet interfaces, the Frame Relay, PPP, or Cisco HDLC headers are counted, while the bit or byte stuffing headers are excluded.

To enable accounting of Layer 2 overhead at the PIC level for ingress and egress traffic on interfaces:

1. Access a DPC or an MPC-occupied slot and the PIC where the interface is to be enabled.

   [edit chassis]
   user@host# **edit fpc *slot-number* pic *number***

2. Specify the Layer 2 overhead value in bytes that is the octet adjustment per packet added to the total octet count for ingress and egress traffic on all the interfaces in the PIC.

   [edit chassis fpc *slot-number* pic *number*]
   user@host# **set account-layer2-overhead**

See Also

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics on page 533
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 537
- **account-layer2-overhead on page 1067**

- *Ethernet Interfaces Feature Guide for Routing Devices*

Verifying the Accounting of Layer 2 Overhead in Interface Statistics

Purpose Display information about the Layer 2 overhead bytes that are counted in interface statistics for egress and ingress traffic on Ethernet interfaces.

Action

- To display information about the Layer 2 overhead bytes that are counted in interface statistics:

  NOTE: For physical and logical interfaces, the values displayed for the Input rate and Output rate fields under the Traffic statistics section include the Layer 2 overhead of the packets.

  user@host> **show interfaces ge-5/2/0 statistics detail**

```
Physical interface: ge-5/2/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 519, Generation: 149
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  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: 00:1d:b5:61:d9:74, Hardware address: 00:1d:b5:61:d9:74
  Last flapped   : 2009-11-11 11:24:00 PST (09:23:08 ago)
  Statistics last cleared: 2009-11-11 17:50:58 PST (02:56:10 ago)
  Traffic statistics:
   Input  bytes  :              271524                    0 bps
   Output bytes  :            37769598                  352 bps
   Input  packets:                3664                    0 pps
   Output packets:              885790                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :           16681118
    Input  packets:                  0
    Output packets:             362633
   Multicast statistics:
    IPV4 multicast statistics:
     Input  bytes  :              112048                    0 bps
     Output bytes  :            20779920                    0 bps
     Input  packets:                1801                    0 pps
     Output packets:              519498                    0 pps
    IPV6 multicast statistics:
     Input  bytes  :              156500                    0 bps
     Output bytes  :            16681118                    0 bps
     Input  packets:                1818                    0 pps
     Output packets:              362633                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2 channel
 errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0, HS link
 CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort               882558               882558                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                3232                 3232                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :              271524
     Output bytes  :            37769598
     Input  packets:                3664
     Output packets:              885790
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :           16681118
      Input  packets:                  0
      Output packets:             362633
    Local statistics:
     Input  bytes  :              271524
     Output bytes  :              308560
     Input  packets:                3664
     Output packets:                3659
    Transit statistics:
```



```
 Input  bytes  :                  0                    0 bps
 Output bytes  :           37461038                    0 bps
 Input  packets:                  0                    0 pps
 Output packets:             882131                    0 pps
 IPv6 transit statistics:
  Input  bytes  :                 0
  Output bytes  :          16681118
  Input  packets:                 0
  Output packets:            362633
Multicast statistics:
 IPV4 multicast statistics:
  Input  bytes  :            112048                    0 bps
  Output bytes  :          20779920                    0 bps
  Input  packets:              1801                    0 pps
  Output packets:            519498                    0 pps
 IPV6 multicast statistics:
  Input  bytes  :            156500                    0 bps
  Output bytes  :          16681118                    0 bps
  Input  packets:              1818                    0 pps
  Output packets:            362633                    0 pps
Protocol inet, MTU: 1500, Generation: 151, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 40.40.40.0/30, Local: 40.40.40.2, Broadcast: 40.40.40.3, Generation: 167
Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: ::40.40.40.0/126, Local: ::40.40.40.2
Generation: 169
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974
Protocol multiservice, MTU: Unlimited, Generation: 171
Generation: 153, Route table: 0
  Policer: Input: __default_arp_policer__
```

Related Documentation

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics on page 533
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 536
- **show interfaces on page 1857**
- *show interfaces statistics*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gigabit Ethernet Policers

Overview

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can define rate limits for premium and aggregate traffic received on the interface. These policers allow you to perform simple traffic policing without configuring a firewall filter. First you configure the Ethernet policer profile, next you classify ingress and egress traffic, then you can apply the policer to a logical interface.

For Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the policer rates you configure can be different than the rates on the Packet Forward Engine. The difference results from Layer 2 overhead. The PIC accounts for this difference.

NOTE:
On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

For information on understanding Ethernet frame statistics, see the *MX Series Layer 2 Configuration Guide*.

Configuring a Policer

To configure an Ethernet policer profile, include the **ethernet-policer-profile** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
ethernet-policer-profile {
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
}
```

In the Ethernet policer profile, the aggregate-priority policer is mandatory; the premium-priority policer is optional.

For aggregate and premium policers, you specify the bandwidth limit in bits per second. You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 4.29 Gbps.

The maximum burst size controls the amount of traffic bursting allowed. To determine the burst-size limit, you can multiply the bandwidth of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur:

burst size = *bandwidth* x *allowable time for burst traffic*

If you do not know the interface bandwidth, you can multiply the maximum MTU of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. The burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 MB.

Specifying an Input Priority Map

An input priority map identifies ingress traffic with specified IEEE 802.1p priority values, and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an input priority map by including the **ieee802.1 premium** statement at the **[edit interfaces *interface-name* gigether-options ethernet-policer-profile input-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]
ieee802.1p premium [ values ];
```

The priority values can be from 0 through 7. The remaining traffic is classified as nonpremium (or aggregate). For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 544.

NOTE: **On IQ2 and IQ2-E interfaces and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.**

Specifying an Output Priority Map

An output priority map identifies egress traffic with specified queue classification and packet loss priority (PLP), and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an output priority map by including the **classifier** statement at the **[edit interfaces *interface-name* gigether-options ethernet-policer-profile output-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile
  output-priority-map]
classifier {
  premium {
```

```
      forwarding-class class-name {
        loss-priority (high | low);
      }
    }
  }
```

You can define a forwarding class, or you can use a predefined forwarding class. Table 52 on page 542 shows the predefined forwarding classes and their associated queue assignments.

Table 52: Default Forwarding Classes

Forwarding Class Name	Queue
best-effort	Queue 0
expedited-forwarding	Queue 1
assured-forwarding	Queue 2
network-control	Queue 3

For more information about CoS forwarding classes, see the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*. For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 544.

Applying a Policer

On all MX Series Router interfaces, Gigabit Ethernet IQ, IQ2, and IQ2-E PICs, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can apply input and output policers that define rate limits for premium and aggregate traffic received on the logical interface. Aggregate policers are supported on Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

These policers allow you to perform simple traffic policing without configuring a firewall filter.

To apply policers to specific source MAC addresses, include the **accept-source-mac** statement:

```
accept-source-mac {
  mac-address mac-address {
    policer {
      input cos-policer-name;
      output cos-policer-name;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]

- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where ***n*** is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include multiple **mac-address** statements in the logical interface configuration.

NOTE: On untagged Gigabit Ethernet interfaces you should not configure the **source-address-filter** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces you should not configure the **source-address-filter** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

NOTE: If the remote Ethernet card is changed, the interface does not accept traffic from the new card because the new card has a different MAC address.

The MAC addresses you include in the configuration are entered into the router's MAC database. To view the router's MAC database, enter the **show interfaces mac-database *interface-name*** command:

user@host> **show interfaces mac-database *interface-name***

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.

NOTE: On IQ2 and IQ2-E PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

You can use the same policer one or more times.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

Configuring MAC Address Filtering

You cannot explicitly define traffic with specific source MAC addresses to be rejected; however, for Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet DPCs on MX Series routers, you can block all incoming packets that do not have a source address specified in the **accept-source-mac** statement. For more information about the **accept-source-mac** statement, see "Applying a Policer" on page 542.

To enable this blocking, include the **source-filtering** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level:

```
[edit interfaces interface-name gigether-options]
source-filtering;
```

For more information about the **source-filtering** statement, see "Configuring MAC Address Filtering for Ethernet Interfaces" on page 14.

To accept traffic even though it does not have a source address specified in the **accept-source-mac** statement, include the **no-source-filtering** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level:

```
[edit interfaces interface-name gigether-options]
no-source-filtering;
```

Example: Configuring Gigabit Ethernet Policers

Example

This example illustrates the following:

- Configure interface **ge-6/0/0** to treat priority values 2 and 3 as premium. On ingress, this means that IEEE 802.1p priority values **2** and **3** are treated as premium. On egress, it means traffic that is classified into queue 0 or 1 with PLP of low and queue 2 or 3 with PLP of high, is treated as premium.
- Define a policer that limits the premium bandwidth to 100 Mbps and burst size to 3 k, and the aggregate bandwidth to 200 Mbps and burst size to 3 k.
- Specify that frames received from the MAC address **00:01:02:03:04:05** and the VLAN ID **600** are subject to the policer on input and output. On input, this means frames received with the source MAC address **00:01:02:03:04:05** and the VLAN ID 600 are subject to the policer. On output, this means frames transmitted from the router with the destination MAC address **00:01:02:03:04:05** and the VLAN ID **600** are subject to the policer.

Example Configuration

```
[edit interfaces]
ge-6/0/0 {
   gigether-options {
      ether-switch-profile {
         ether-policer-profile {
            input-priority-map {
               ieee-802.1p {
                  premium [ 2 3 ];
               }
            }
            output-priority-map {
               classifier {
                  premium {
                     forwarding-class best-effort {
                        loss-priority low;
                     }
                     forwarding-class expedited-forwarding {
                        loss-priority low;
                     }
                     forwarding-class assured-forwarding {
                        loss-priority high;
                     }
                     forwarding-class network-control {
                        loss-priority high;
                     }
                  }
               }
            }
            policer policer-1 {
               premium {
                  bandwidth-limit 100m;
                  burst-size-limit 3k;
               }
               aggregate {
                  bandwidth-limit 200m;
                  burst-size-limit 3k;
               }
            }
         }
      }
   }
   unit 0 {
      accept-source-mac {
         mac-address 00:01:02:03:04:05 {
            policer {
               input policer-1;
               output policer-1;
            }
         }
      }
   }
}
```


Related Documentation

- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 529
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 546
- Configuring MAC Address Accounting on page 531
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

Overview

For Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers, you can configure two-color and tricolor marking policers and apply them to logical interfaces to prevent traffic on the interface from consuming bandwidth inappropriately.

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks router architectures support three types of policer:

- Two-color policer—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit in some way, or simply discard them. A policer is most useful for metering traffic at the port (physical interface) level.
- Single-rate tricolor marking (single-rate TCM)—A single-rate tricolor marking policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).

  Starting in Junos OS Release 13.1, traffic is classified into three categories: Green, Red, and Yellow. Following list describes the categories:

  - Green—Burst size of the packets that arrive is less than the sum of the configured CIR and CBS.

  - Red—Burst size of the packets that arrive is greater than the sum of the configured CIR and EBS.
  - Yellow—Burst size of the packets that arrive is greater than the CBS but less than the EBS.

  Single-rate TCM is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate Tricolor Marking (two-rate TCM)—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and peak information rate (PIR), along with their associated burst sizes, the CBS and EBS.

  Traffic is classified into the following three categories:

  - Green—Burst size of the packets that arrive is less than the sum of the configured CIR and CBS.
  - Red—Burst size of the packets that arrive is greater than the sum of the configured PIR and EBS.
  - Yellow—Traffic does not belong to either the green or the red category.

  Two-rate TCM is most useful when a service is structured according to arrival rates and not necessarily packet length.

NOTE: **Unlike policing (described in** "Configuring Gigabit Ethernet Policers" on page 539**), configuring two-color policers and tricolor marking policers requires that you configure a firewall filter.**

Configuring a Policer

Two-color and tricolor marking policers are configured at the **[edit firewall]** hierarchy level.

A tricolor marking policer polices traffic on the basis of metering rates, including the CIR, the PIR, their associated burst sizes, and any policing actions configured for the traffic.

To configure tricolor policer marking, include the **three-color-policer** statement with options at the **[edit firewall]** hierarchy level:

```
[edit firewall]
three-color-policer name {
  action {
    loss-priority high {
      then discard;
    }
  }
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
```

```
      committed-burst-size bytes;
      excess-burst-size bytes;
    }
    two-rate {
      (color-aware | color-blind);
      committed-information-rate bps;
      committed-burst-size bytes;
      peak-information-rate bps;
      peak-burst-size bytes;
    }
  }
```

For more information about configuring tricolor policer markings, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide* and the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*.

Applying a Policer

Apply a two-color policer or tricolor policer to a logical interface to prevent traffic on the interface from consuming bandwidth inappropriately. To apply two-color or tricolor policers, include the **layer2-policer** statement:

```
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  policer-name;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Use the **input-policer** statement to apply a two-color policer to received packets on a logical interface and the **input-three-color** statement to apply a tricolor policer. Use the **output-policer** statement to apply a two-color policer to transmitted packets on a logical interface and the **output-three-color** statement to apply a tricolor policer. The specified policers must be configured at the **[edit firewall]** hierarchy level. For each interface, you can configure a three-color policer or two-color input policer or output policers—you cannot configure both a three-color policer and a two-color policer.

Example: Configuring and Applying a Policer

Configure tricolor policers and apply them to an interface:

```
[edit firewall]
three-color-policer three-color-policer-color-blind {
  logical-interface-policer;
  two-rate {
    color-blind;
    committed-information-rate 1500000;
    committed-burst-size 150;
    peak-information-rate 3;
```



```
    peak-burst-size 300;
  }
}
three-color-policer three-color-policer-color-aware {
  logical-interface-policer;
  two-rate {
    color-aware;
    committed-information-rate 1500000;
    committed-burst-size 150;
    peak-information-rate 3;
    peak-burst-size 300;
  }
}
[edit interfaces ge-1/1/0]
unit 1 {
  layer2-policer {
    input-three-color three-color-policer-color-blind;
    output-three-color three-color-policer-color-aware;
  }
}
```

Configure a two-color policer and apply it to an interface:

```
[edit firewall]
policer two-color-policer {
  logical-interface-policer;
  if-exceeding {
    bandwidth-percent 90;
    burst-size-limit 300;
  }
  then loss-priority-high;
}
[edit interfaces ge-1/1/0]
unit 2 {
  layer2-policer {
    input-policer two-color-policer;
    output-policer two-color-policer;
  }
}
```

Release History Table

Release	Description
13.1	Starting in Junos OS Release 13.1, traffic is classified into three categories: Green, Red, and Yellow.

Related Documentation

- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 529
- Configuring Gigabit Ethernet Policers on page 539
- Configuring MAC Address Accounting on page 531
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 29

Configuring Gigabit Ethernet Autonegotiation

Gigabit Ethernet Autonegotiation Overview

Autonegotiation is enabled by default on all Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. However, you can explicitly enable autonegotiation to configure remote fault options manually.

NOTE:

- When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
- On ACX Series Universal Metro Routers, when the autonegotiation is disabled, the speed has to be explicitly configured to 10–100 Mbps.
- On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

Related Documentation

- Configuring Gigabit Ethernet Autonegotiation on page 551

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gigabit Ethernet Autonegotiation

Configuring Gigabit Ethernet Autonegotiation with Remote Fault

To configure explicit autonegotiation and remote fault, include the **auto-negotiation** statement and the **remote-fault** option at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level.

```
[edit interfaces ge-fpc/pic/port gigether-options]
(auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
  local-interface-offline>
```

Configuring Flow Control

To enable flow control, include the **flow-control** statement at the **[edit interfaces ge-*fpc/pic>/port* gigether-options]** hierarchy level. For more information, see "Configuring Flow Control" on page 12.

Configuring Autonegotiation Speed on MX Series Routers

MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs support autonegotiation of speed. The autonegotiation specified interface speed is propagated to CoS, routing protocols, and other system components. Half-duplex mode is not supported.

MX Series routers with IQ2 PICs connected to other devices require matching auto-negotiation configurations for both the PIC and for the device in order to achieve link up.

To specify the autonegotiation speed, use the **speed (auto | 1Gbps | 100Mbps | 10Mbps | auto-10m-100m)** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level.

To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**. If the negotiated speed and the interface speed do not match, the link will not be brought up.

If you set the autonegotiation speed **auto** option, then the port speed is negotiated.

Starting from Junos OS Release 14.2, the**auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m**maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.

You can disable auto MDI/MDIX using the **no-auto-mdix** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level.

Use the **show interfaces ge-*fpc/pic/port* brief** command to display the auto negotiation of speed and auto MDI/MDIX states.

NOTE: Starting in Junos OS Release 14.2, on MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the **auto-10m-100m** option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the **show interfaces** command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the **show interfaces** command only. The actual interface speed is 100 Mbps.

Displaying Autonegotiation Status

To display Gigabit Ethernet interface details, including the autonegotiation status, use the operational mode command **show interfaces ge-** ***fpc/pic/port*** **extensive**.

Table 53 on page 553 and Table 54 on page 555 provide information about the autonegotiation status on local and remote routers with fiber interfaces. The status of the link and LED can vary depending on the level of autonegotiation set and the transmit and receive fiber status.

Table 53: Mode and Autonegotiation Status (Local)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Red	DOWN	
ON	ON	Default	Green	UP	No-autonegotiation
ON	OFF	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Green	UP	
ON	ON	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	OFF	No-autonegotiation	Red	DOWN	

Table 53: Mode and Autonegotiation Status (Local) (continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit	Green	UP	No-autonegotiation
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Green	UP	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Green	UP	Complete
OFF	ON	Explicit+RFI-Offline	Red	DOWN	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Green	UP	No-autonegotiation
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Offline	Green	UP	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Offline	Red	DOWN	

Table 53: Mode and Autonegotiation Status (Local) (continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	No-autonegotiation*
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Green	UP	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	Complete

Table 54: Mode and Autonegotiation Status (Remote)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	ON	Default	Red	DOWN	
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	ON	No-autonegotiation	Red	DOWN	

Table 54: Mode and Autonegotiation Status (Remote) (continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	ON	Explicit	Red	DOWN	
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Offline	Red	DOWN	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	Complete
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	

Release History Table

Release	Description
14.2	Starting from Junos OS Release 14.2, the**auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m**maximum speed.
14.2	Starting in Junos OS Release 14.2, on MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the **auto-10m-100m** option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link.

Related Documentation

- Gigabit Ethernet Autonegotiation Overview on page 551
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 30

Stacking and Rewriting Gigabit Ethernet VLAN Tags

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

Stacking and rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between customer edge (CE) routers that share one VLAN ID. A frame can be received on an interface, or it can be internal to the system (as a result of the **input-vlan-map** statement).

On IQ2 interfaces, 10-Gigabit Ethernet LAN/WAN PIC, 40-Gigabit Ethernet MIC, 100-Gigabit Ethernet MIC, IQ2-E interfaces, and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

NOTE: When swap-by-poppush is configured on the interface, when a VLAN tag is swapped, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN being swapped. If swap-by-poppush is not configured on the interface, the VLAN IEEE 802.1p bits of the of the VLAN being swapped remains same.

You can stack and rewrite VLAN tags on the following interfaces:

- Gigabit Ethernet
- Gigabit Ethernet IQ
- 10-Gigabit Ethernet LAN/WAN PIC
- 40-Gigabit Ethernet MIC
- 100-Gigabit Ethernet MIC
- Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, and MX Series router Gigabit Ethernet Interfaces
- Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with the VLAN encapsulation type configured to support Layer 2 tunneling protocols such as circuit cross-connect (CCC) or virtual private LAN service (VPLS) (as described in “802.1Q VLANs Overview” on page 244)

Related Documentation

- 802.1Q VLANs Overview on page 244
- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 560
- *Ethernet Interfaces Feature Guide for Routing Devices*

Stacking and Rewriting Gigabit Ethernet VLAN Tags

You can configure rewrite operations to stack (**push**), remove (**pop**), or rewrite (**swap**) tags on single-tagged frames and dual-tagged frames. If a port is not tagged, rewrite operations are not supported on any logical interface on that port.

You can configure the following VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **pop-pop**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.
- **push**—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- **push-push**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, push two VLAN tags in front of the frame.
- **swap-push**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, replace the outer VLAN tag of the frame with 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**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value.

You configure VLAN rewrite operations for logical interfaces in the input VLAN map for incoming frames and in the output VLAN map for outgoing frames. To configure the input VLAN map, include the **input-vlan-map** statement:

```
input-vlan-map {
  ...interface-specific configuration...
}
```

To configure the output VLAN map, include the **output-vlan-map** statement:

```
output-vlan-map {
  ...interface-specific configuration...
}
```

You can include both statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The type of VLAN rewrite operation permitted depends upon whether the frame is single-tagged or dual-tagged. Table 55 on page 562 shows supported rewrite operations and whether they can be applied to single-tagged frames or dual-tagged frames. The table also indicates the number of tags being added or removed during the operation.

Table 55: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames

Rewrite Operation	Untagged	Single-Tagged	Dual-Tagged	Number of Tags
pop	No	Yes	Yes	– 1
push	Sometimes	Yes	Yes	+1
swap	No	Yes	Yes	0
push-push	Sometimes	Yes	Yes	+2
swap-push	No	Yes	Yes	+1
swap-swap	No	No	Yes	0
pop-pop	No	No	Yes	– 2
pop-swap	No	No	Yes	– 1

The rewrite operations **push** and **push-push** can be valid in certain circumstances on frames that are not tagged. For example, a single-tagged logical interface (interface 1) and a dual-tagged logical interface (interface 2) have the following configurations:

Interface 1

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop;
}
output-vlan-map {
  push;
}
```

Interface 2

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-pop;
}
output-vlan-map {
  push-push;
}
```

When a frame is received on the interface as a result of the **input-vlan-map** operation, the frame is not tagged. As it goes out of the second interface, the **output-vlan-map** operation **push-push** is applied to it. The resulting frame will be dual-tagged at the logical interface output.

Depending on the VLAN rewrite operation, you configure the rewrite operation for the interface in the input VLAN map, the output VLAN map, or in both the input VLAN map and the output VLAN map. Table 56 on page 563 shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

Table 56: Applying Rewrite Operations to VLAN Maps

	Output VLAN Map								
Input VLAN Map	none	push	pop	swap	push-push	swap-push	swap-swap	pop-pop	swap-pop
none	Yes	No	No	Yes	No	No	Yes	No	No
push	No	No	Yes	No	No	No	No	No	No
pop	No	Yes	No	No	No	No	No	No	No
swap	Yes	No	No	Yes	No	No	No	No	No
push-push	No	No	No	No	No	No	No	Yes	No
swap-push	No	No	No	No	No	No	No	No	Yes
swap-swap	Yes	No	No	No	No	No	Yes	No	No
pop-pop	No	No	No	No	Yes	No	No	No	No
pop-swap	No	No	No	No	No	Yes	No	No	No

You must know whether the VLAN rewrite operation is valid and is applied to the input VLAN map or the output VLAN map. You must also know whether the rewrite operation requires you to include statements to configure the inner and outer TPIDs and inner and outer VLAN IDs in the input VLAN map or output VLAN map. For information about configuring inner and outer TPIDs and inner and outer VLAN IDs, see "Configuring Inner and Outer TPIDs and VLAN IDs" on page 566.

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559
- Understanding swap-by-poppush on page 587
- **swap-by-poppush on page 1379**

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames

For Gigabit Ethernet IQ interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure frames with particular TPIDs to be processed as tagged frames. To do this, you specify up to eight IEEE 802.1Q TPID values per port; a frame with any of the specified TPIDs is processed as a tagged frame; however, with IQ2 and IQ2-E interfaces, only the first four IEEE 802.1Q TPID values per port are supported. To configure the TPID values, include the **tag-protocol-id** statement:

tag-protocol-id [*tpids*];

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]**
- **[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]**

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** or **[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559
- **aggregated-ether-options on page 1079**
- **ethernet-switch-profile on page 1154**
- **gigether-options on page 1180**
- **tag-protocol-id on page 1387**
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Tag Protocol IDs (TPIDs) on PTX Series Packet Transport Routers

This topic describes how to configure the TPIDs expected to be sent or received on a particular VLAN for PTX Series Packet Transport Routers.

For other types of Juniper Networks Ethernet PICs, you could configure 8 TPIDs per port. However, the PTX Series Packet Transport Routers use MTIP and TL to classify a specific TPID and Ethernet type. For MTIP, you can configure a maximum of 8 TPIDs for each MAC chip.

As a consequence, you can specify the **tag-protocol-id** configuration statement only for the first port (**0**) of a PTX Series Ethernet PIC. If you configure **tag-protocol-id** statements on the other port, the configuration is ignored and a system error is recorded.

For example, the following is a supported configuration:

```
[edit interfaces et-2/0/0]
gigether-options {
  ethernet-switch-profile {
    tag-protocol-id [0x8100 0x9100];
  }
}
```

The **tag-protocol-id** configuration statement supports up to eightTPIDs on port **0** of a given Ethernet PIC. All eight TPIDs are populated to the two MTIPs and TLs associated with the Ethernet PIC.

Related Documentation

- Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames on page 563
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 249

Configuring Stacked VLAN Tagging

To configure stacked VLAN tagging for all logical interfaces on a physical interface:

1. In configuration mode, go to the **[edit interfaces interface-name]** hierarchy level.

   [edit]
   user@host# edit interfaces *interface-name*

2. Include the **stacked-vlan-tagging** statement.

   [edit interfaces *interface-name*]
   user@host# set stacked-vlan-tagging

If you include the **stacked-vlan-tagging** statement in the configuration, you must configure dual VLAN tags for all logical interfaces on the physical interface. For more information, see "Stacking a VLAN Tag" on page 569.

Related Documentation

- **stacked-vlan-tagging on page 1372**
- Stacking a VLAN Tag on page 569
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Dual VLAN Tags

To configure dual VLAN tags on a logical interface, include the **vlan-tags** statement:

vlan-tags inner <*tpid*.>*vlan-id* outer <*tpid*.>*vlan-id*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The outer tag VLAN ID range is from 1 through 511 for normal interfaces, and from 512 through 4094 for VLAN CCC or VLAN VPLS interfaces. The inner tag is not restricted.

You must also include the **stacked-vlan-tagging** statement in the configuration. See "Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags" on page 579.

Related Documentation

- **unit on page 1420**
- Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 579
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Inner and Outer TPIDs and VLAN IDs

For some rewrite operations, you must configure the inner or outer tag-protocol identifier (TPID) values and inner or outer virtual local area network identifier (VLAN ID) values. These values can be applied to either the input VLAN map or the output VLAN map.

1. On Ethernet IQ, IQ2, and IQ2-E interfaces; on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces; and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to configure the inner TPID, include the **inner-tag-protocol-id** statement at the following hierarchy levels:

   - **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
   - **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
   - **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
   - **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

   ```
   user@host# set inner-tag-protocol-id tpid;
   ```

2. For the inner VLAN ID, include the **inner-vlan-id** statement. For the outer TPID, include the **tag-protocol-id** statement. For the outer VLAN ID, include the **vlan-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

   ```
   input-vlan-map {
     (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
     inner-tag-protocol-id tpid;
     inner-vlan-id number;
     tag-protocol-id tpid;
     vlan-id number;
   }
   output-vlan-map {
     (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
     inner-tag-protocol-id tpid;
     inner-vlan-id number;
     tag-protocol-id tpid;
     vlan-id number;
   }
   ```

3. For aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, include the **tag-protocol-id** statement for the outer TPID. For the outer VLAN ID, include the **vlan-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

   ```
   input-vlan-map {
   ```



```
        (pop | push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    output-vlan-map {
        (pop | push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
```

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see "802.1Q VLANs Overview" on page 244.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level or **[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level. For more information, see "Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames" on page 563.

Table 57 on page 567 and Table 58 on page 568 specify when these statements are required. Table 57 on page 567 indicates valid statement combinations for rewrite operations for the input VLAN map. "No" means the statement must not be included in the input VLAN map for the rewrite operation. "Optional" means the statement may be optionally specified for the rewrite operation in the input VLAN map. "Any" means that you must include the **vlan-id** statement, **tag-protocol-id** statement, **inner-vlan-id** statement, or **inner-tag-protocol-id** statement.

Table 57: Rewrite Operations and Statement Usage for Input VLAN Maps

	Input VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Optional	Optional	No	No
pop	No	No	No	No
swap	Any	Any	No	No
push-push	Optional	Optional	Optional	optional
swap-push	Optional	Optional	Any	Any
swap-swap	Optional	Optional	Any	Any
pop-swap	No	No	Any	Any
pop-pop	No	No	No	No

Table 58 on page 568 indicates valid statement combinations for rewrite operations for the output VLAN map. “No” means the statement must not be included in the output VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 58: Rewrite Operations and Statement Usage for Output VLAN Maps

	Output VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	No	Optional	No	No
pop	No	No	No	No
swap	No	Optional	No	No
push-push	No	Optional	No	Optional
swap-push	No	Optional	No	Optional
swap-swap	No	Optional	No	Optional
pop-swap	No	No	No	Optional
pop-pop	No	No	No	No

Input VLAN Map with inner-vlan-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
  inner-tag-protocol-id tpid;
```



```
}
```

Related Documentation

- inner-tag-protocol-id on page 1197
- input-vlan-map on page 1203
- output-vlan-map on page 1295
- pop-swap on page 1308
- unit on page 1420
- *Ethernet Interfaces*

Stacking a VLAN Tag

To stack a VLAN tag on all tagged frames entering or exiting the interface, include the **push**, **vlan-id**, and **tag-protocol-id** statements in the input VLAN map or the output VLAN map:

```
input-vlan-map {
  push;
  vlan-id number;
  tag-protocol-id tpid;
}
output-vlan-map {
  push;
  tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

If you include the **push** statement in an interface's input VLAN map, see Table 56 on page 563 for information about permissible rewrite operations,

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see "802.1Q VLANs Overview" on page 244.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level. For more information, see "Configuring Inner and Outer TPIDs and VLAN IDs" on page 566.

Related Documentation

- **tag-protocol-id on page 1388**
- **unit on page 1420**
- Table 56 on page 563
- 802.1Q VLANs Overview on page 244
- Configuring Inner and Outer TPIDs and VLAN IDs on page 566
- *Ethernet Interfaces Feature Guide for Routing Devices*

Stacking Two VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to push two VLAN tags in front of tagged frames entering or exiting the interface, include the **push-push** statement in the input VLAN map or the output VLAN map:

```
push-push;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- **input-vlan-map on page 1203**
- **output-vlan-map on page 1295**
- **pop on page 1306**
- **unit on page 1420**
- See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- *Ethernet Interfaces Feature Guide for Routing Devices*

Removing a VLAN Tag

To remove a VLAN tag from all tagged frames entering or exiting the interface, include the **pop** statement in the input VLAN map or output VLAN map:

```
pop;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295
- pop on page 1306
- unit on page 1420
- *Ethernet Interfaces Feature Guide for Routing Devices*

Removing the Outer and Inner VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to remove both the outer and inner VLAN tags of the frame, include the **pop-pop** statement in the input VLAN map or output VLAN map:

```
pop-pop;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295
- pop-pop on page 1307
- unit on page 1420

- See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- *Ethernet Interfaces Feature Guide for Routing Devices*

Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame with a user-specified VLAN tag value, include the **pop-swap** statement in the input VLAN map or output VLAN map:

```
pop-swap;
```

The inner tag becomes the outer tag in the final frame.

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- **input-vlan-map on page 1203**
- **output-vlan-map on page 1295**
- **pop-swap on page 1308**
- **unit on page 1420**
- See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- *Ethernet Interfaces Feature Guide for Routing Devices*

Rewriting the VLAN Tag on Tagged Frames

To rewrite the VLAN tag on all tagged frames entering the interface to a specified VLAN ID and TPID, include the **swap**, **tag-protocol-id**, and **vlan-id** statements in the input VLAN map:


```
input-vlan-map {
  swap;
  vlan-id number;
  tag-protocol-id tpid;
}
```

To rewrite the VLAN tag on all tagged frames exiting the interface to a specified VLAN ID and TPID, include the **swap** and **tag-protocol-id** statements in the output VLAN map:

```
output-vlan-map {
  swap;
  vlan-id number;
  tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**

You cannot include both the **swap** statement and the **vlan-id** statement in the output VLAN map configuration. If you include the **swap** statement in the configuration, the VLAN ID in outgoing frames is rewritten to the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see "802.1Q VLANs Overview" on page 244.

The swap operation works on the outer tag only, whether or not you include the **stacked-vlan-tagging** statement in the configuration. For more information, see "Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags" on page 579.

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295
- swap on page 1379
- vlan-id on page 1433
- tag-protocol-id on page 1388
- unit on page 1420
- For more information about binding a VLAN ID to the logical interface, see 802.1Q VLANs Overview on page 244.
- For more information about the swap operation, see Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 579.
- *Ethernet Interfaces Feature Guide for Routing Devices*

Rewriting a VLAN Tag on Untagged Frames

Overview

You can rewrite VLAN tags on untagged incoming and outgoing frames with the ethernet-ccc and the ethernet-vpls encapsulations for the following routers:

- M120 routers and M320 routers with:
  - Gigabit Ethernet IQ PIC with SFP
  - Gigabit Ethernet IQ2 PICs with SFP
  - Gigabit Ethernet Enhanced IQ2 (IQ2E) PICs with SFP
  - 10-Gigabit Ethernet IQ2 PIC with XFP
  - 10-Gigabit Ethernet Enhanced IQ2 (IQ2E) PIC with XFP
- MX240, MX480, and MX960 routers with:
  - Gigabit Ethernet Enhanced DPC with SFP
  - Gigabit Ethernet Enhanced Queuing IP Services DPCs with SFP
  - 10-Gigabit Ethernet Enhanced DPCs with XFP
  - 10-Gigabit Ethernet Enhanced Queuing IP Services DPC with XFP

On M Series routers with Gigabit Ethernet IQ2 PICs and Gigabit Ethernet Enhanced IQ2 (IQ2E) PICs, you can perform all the rewrite VLAN tag operations.

Consider a network where two provider edges (PE) are connected by a Layer 2 circuit. PE1 is receiving traffic on an untagged port while the corresponding port on PE2 is tagged. In the normal case, packets coming from PE1 will be dropped at PE2 because it is expecting tagged packets. However, if PE1 can push a VLAN tag on the incoming packet before sending it across to PE2, you can ensure that packets are not dropped. To make it work in both directions, PE1 must strip the VLAN tag from outgoing packets. Therefore, a push on the ingress side is always paired with a pop on the egress side.

The rewrite operations represented by the following statement options are supported under **ethernet-ccc** and **ethernet-vpls** encapsulations:

- **push**—A VLAN tag is added to the incoming untagged frame.
- **pop**—VLAN tag is removed from the outgoing frame.

- **push-push**—An outer and inner VLAN tag are added to the incoming untagged frame.
- **pop-pop**—Both the outer and inner VLAN tags of the outgoing frame are removed.

IQ2 and 10-Gigabit Ethernet PICs support all rewrite operations described above. Details on the possible combinations of usage are explained later in this section.

NOTE: The **push-push** and **pop-pop** operations are not supported on the Gigabit Ethernet IQ PIC.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the tag parameters have to be explicitly specified. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames. Table 59 on page 575 through Table 61 on page 575 explain the rules in more detail.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the **vlan-id** parameters (**vlan-id** for **push** and **vlan-id** or **inner-vlan-id** for **push-push**) have to be explicitly specified. TPID however, is optional and the default value of **0x8100** is set if not configured. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames.

Table 59: Input VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Yes	Optional	No	Optional
push-push	Yes	Optional	Yes	Optional

Table 60: Output VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
pop	No	No	No	No
pop-pop	No	No	No	No

Table 61: Rules for Applying Rewrite Operations to VLAN Maps

	Output VLAN Map		
Input VLAN Map	None	pop	pop-pop
None	Yes	No	No
push	No	Yes	No

Table 61: Rules for Applying Rewrite Operations to VLAN Maps (continued)

push-push	No	No	Yes

You can use the **show interface *interface-name*** command to display the status of a modified VLAN map for the specified interface.

Example: push and pop with Ethernet CCC Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 600;
    }
    output-vlan-map pop;
    family ccc;
  }
}
```

Example: push-push and pop-pop with Ethernet CCC Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family ccc;
  }
}
```

Example: push and pop with Ethernet VPLS Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 700;
    }
    output-vlan-map pop;
    family vpls;
  }
}
```

Example: push-push and pop-pop with Ethernet VPLS Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family vpls;
  }
}
```

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295
- pop on page 1306
- pop-pop on page 1307
- push on page 1327
- push-push on page 1328
- unit on page 1420

- *Ethernet Interfaces Feature Guide for Routing Devices*

Rewriting a VLAN Tag and Adding a New Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and on Gigabit Ethernet and 10-Gigabit Ethernet interfaces on EX Series switches, to replace the outer VLAN tag of the incoming frame with a user-specified VLAN tag value, include the **swap-push** statement in the input VLAN map or output VLAN map:

```
swap-push
```

A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295
- swap-push on page 1380
- unit on page 1420
- *Ethernet Interfaces Feature Guide for Routing Devices*

Rewriting the Inner and Outer VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value, include the **swap-swap** statement in the input VLAN map or output VLAN map:

```
swap-swap;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 57 on page 567 and Table 58 on page 568 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295
- swap-swap on page 1381
- unit on page 1420
- *Ethernet Interfaces Feature Guide for Routing Devices*

Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags

Configure a VLAN CCC tunnel in which Ethernet frames enter the tunnel at interface **ge-4/0/0** and exit the tunnel at interface **ge-4/2/0**.

The following examples show how to perform the following tasks:

Push a TPID and VLAN ID Pair on Ingress

```
[edit interfaces]
ge-4/0/0 {
   vlan-tagging;
   encapsulation vlan-ccc;
   gigether-options {
      ethernet-switch-profile {
         tag-protocol-id 0x9909;
      }
   }
   unit 0 {
      encapsulation vlan-ccc;
      vlan-id 512;
      input-vlan-map {
         push;
         tag-protocol-id 0x9909;
         vlan-id 520;
      }
      output-vlan-map pop;
   }
}
ge-4/2/0 {
   vlan-tagging;
   encapsulation vlan-ccc;
   unit 0 {
      encapsulation vlan-ccc;
```



```
        vlan-id 515;
        input-vlan-map {
            swap-push;
            vlan-id 520;
            inner-vlan-id 512;
        }
        output-vlan-map {
            pop-swap;
        }
    }
    [edit protocols]
    mpls {
        interface ge-4/0/0.0;
        interface ge-4/2/0.0;
    }
    connections {
        interface-switch vlan-tag-push {
            interface ge-4/0/0.0;
            interface ge-4/2/0.0;
        }
    }
}
```

Stack Inner and Outer VLAN Tags

```
[edit interfaces]
ge-0/2/0 {
    stacked-vlan-tagging;
    mac 00.01.02.03.04.05;
    gigether-options {
        loopback;
    }
    unit 0 {
        vlan-tags outer 0x8100.200 inner 0x8100.200;
    }
}
```

Swap a VLAN ID on Ingress

```
[edit interfaces]
ge-4/0/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id 0x9100;
        }
    }
    ...
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 1000;
        input-vlan-map {
            swap;
            tag-protocol-id 0x9100;
            vlan-id 2000;
        }
    }
```



```
}
ge-4/2/0 {
   vlan-tagging;
   encapsulation vlan-ccc;
   ...
   unit 1 {
      encapsulation vlan-ccc;
      vlan-id 2000;
      input-vlan-map {
         swap;
         tag-protocol-id 0x9100;
         vlan-id 1000;
      }
   }
   [edit protocols]
   mpls {
      ...
      interface ge-4/0/0.1;
      interface ge-4/2/0.1;
   }
   connections {
      ...
      interface-switch vlan-tag-swap {
         interface ge-4/2/0.1;
         interface ge-4/0/0.1;
      }
   }
}
```

Swap a VLAN ID on Egress

```
[edit interfaces]
ge-4/0/0 {
   vlan-tagging;
   encapsulation vlan-ccc;
   ...
   unit 1 {
      encapsulation vlan-ccc;
      vlan-id 1000;
   }
}
ge-4/2/0 {
   vlan-tagging;
   encapsulation vlan-ccc;
   gigether-options {
      ethernet-switch-profile {
         tag-protocol-id 0x8800;
      }
   }
   ...
   unit 1 {
      encapsulation vlan-ccc;
      vlan-id 2000;
      output-vlan-map {
         swap;
         tag-protocol-id 0x8800;
      }
```

```
    }
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}
```

Swap a VLAN ID on Both Ingress and Egress

```
[edit interfaces]
ge-4/0/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
    ...
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 1000;
        input-vlan-map {
            swap;
            tag-protocol-id 0x9100;
            vlan-id 2000;
        }
    }
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        output-vlan-map {
            swap;
            tag-protocol-id 0x8800;
        }
    }
}
[edit protocols]
mpls {
```



```
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}
```

Swap the Outer VLAN Tag and Push a New VLAN Tag on Ingress; Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Egress

```
[edit interfaces]
ge-1/1/0 {
    unit 1 {
        vlan-id 200;
        input-vlan-map {
            swap-push;
            tag-protocol-id 0x9100;
            vlan-id 400;
            inner-tag-protocol-id 0x9100;
            inner-vlan-id 500;
        }
        output-vlan-map {
            pop-swap;
            inner-tag-protocol-id 0x9100;
        }
    }
}
```

Swap a TPID and VLAN ID Pair for Both VLAN Tags on Ingress and on Egress

```
[edit interfaces]
ge-1/1/0 {
    unit 0 {
        vlan-tags {
            inner 0x9100.425;
            outer 0x9200.525;
        }
        input-vlan-map {
            swap-swap;
            tag-protocol-id 0x9100;
            vlan-id 400;
            inner-tag-protocol-id 0x9100;
            inner-vlan-id 500;
        }
        output-vlan-map {
            swap-swap;
            tag-protocol-id 0x9200;
            inner-tag-protocol-id 0x9100;
        }
    }
}
```

Pop the Outer VLAN Tag and Swap the

```
[edit interfaces]
ge-1/1/0 {
```

Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress

```
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map {
      pop-swap;
      tag-protocol-id 0x9100;
      vlan-id 400;
    }
    output-vlan-map{
      swap-push;
      tag-protocol-id 0x9200;
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

Pop a TPID and VLAN ID Pair on Ingress; Push a VLAN ID and TPID Pair on Egress

```
[edit interfaces]
ge-1/1/0 {
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map{
      pop-pop;
    }
    output-vlan-map {
      push-push;
      tag-protocol-id 0x9200;
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

Pop an Outer VLAN Tag to Connect an Untagged VPLS Interface to Tagged VPLS Interfaces

```
[edit interfaces]
ge-1/1/0 {
  vlan-tagging;
  encapsulation extended-vlan-vpls;
  unit 0 {
    vlan-id 0;
    input-vlan-map {
      push;
      vlan-id 0;
    }
    output-vlan-map pop;
    family vpls;
  }
}
```

Related Documentation

- input-vlan-map on page 1203
- output-vlan-map on page 1295

- inner-tag-protocol-id on page 1197
- inner-vlan-id on page 1198
- pop on page 1306
- pop-pop on page 1307
- pop-swap on page 1308
- push on page 1327
- push-push on page 1328
- swap on page 1379
- swap-push on page 1380
- swap-swap on page 1381
- unit on page 1420
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding Transparent Tag Operations and IEEE 802.1p Inheritance

When **swap-by-poppush** is configured on IQ2 interfaces, 10-Gigabit Ethernet LAN/WAN PIC, IQ2-E interfaces, and MX Series interfaces, during a swap operation, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the tag being swapped. If swap-by-poppush is not configured on the interface, the VLAN IEEE 802.1p bits of the tag being swapped remains same.

When **swap-by-poppush** is configured but the incoming packet has no inner VLAN tag (transparent tag), the IEEE 802.1p bits are set to zero .

Table 62 on page 585 describes the relationship between the VLAN map operation and the inheritance of IEEE 802.1p from the transparent tag. It assumes the presence of the transparent tag in the incoming packet. If the transparent tag is not present, the IEEE 802.1p value is set to 0.

Table 62: VLAN Map Operation and IEEE 802.1p Inheritance

Rewrite Operation	Untagged Logical Interface	Transparent tag IEEE 802.1p Inheritance	Single-tagged Logical Interface	Transparent tag IEEE 802.1p Inheritance	Change in number of tags
push-push	yes	OUTER, INNER	NA	no operation	+2
swap-push	NA	no operation	yes	OUTER, INNER	+1
push	yes	OUTER	yes	*none	+1
swap	NA	NA	yes	OUTER	0

Table 62: VLAN Map Operation and IEEE 802.1p Inheritance (continued)

Rewrite Operation	Untagged Logical Interface	Transparent tag IEEE 802.1p Inheritance	Single-tagged Logical Interface	Transparent tag IEEE 802.1p Inheritance	Change in number of tags
NOTE: *In a **push** operation on a single-tagged logical interface, none of the tags (inner, or outer) inherit the IEEE 802.1p bits from the transparent tag.					

The following section shows four different examples of the inheritance of the transparent IEEE 802.1p values into the outer and inner VLAN tags.

Figure 33 on page 586 shows an incoming packet with a transparent tag. A swap-push operation swaps the outer VLAN tag and pushes another VLAN tag. The IEEE 802.1p values are inherited from the transparent tag.

Figure 33: swap-push (transparent tag)

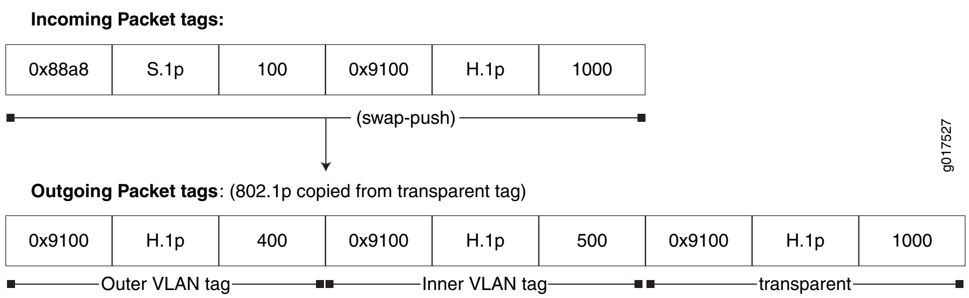

Figure 34 on page 586 shows an incoming packet with no transparent tag. A swap-push operation swaps the outer VLAN tag and pushes another VLAN tag. The IEEE 802.1p value is set to zero, as there is no transparent tag.

Figure 34: swap-push (no transparent tag)

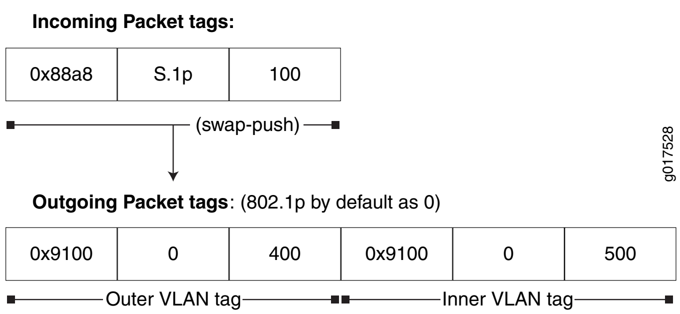

Figure 35 on page 587 shows an incoming packet with a transparent tag. A push operation pushes another VLAN tag. The IEEE 802.1p value is inherited from the transparent tag.

Figure 35: push (transparent tag)

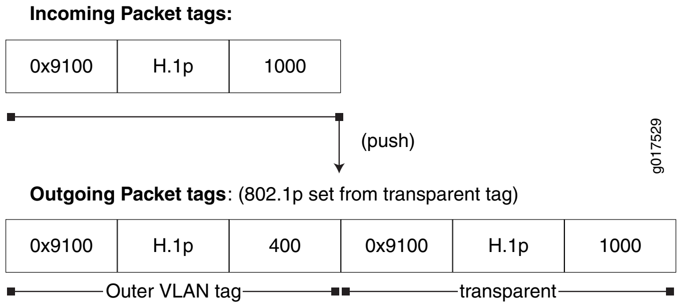

Figure 36 on page 587 shows an incoming packet with a transparent tag. A push-push operation pushes the outer and inner VLAN tags, respectively. The IEEE 802.1p values are inherited from the transparent tag.

Figure 36: push-push (transparent tag)

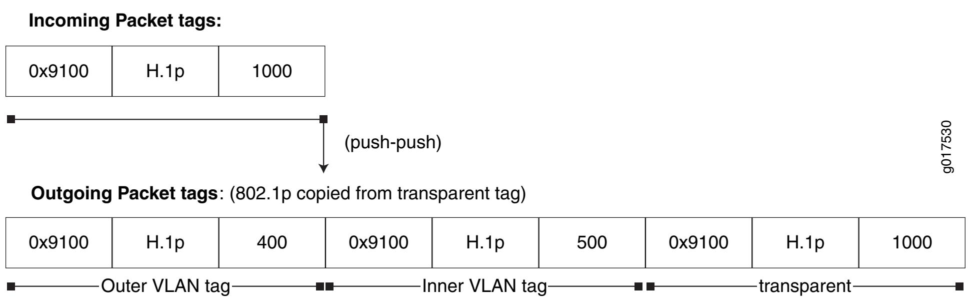

Related Documentation

- *Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag*
- Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag on page 588
- Understanding swap-by-poppush on page 587
- **swap-by-poppush on page 1379**
- *transparent*

Understanding swap-by-poppush

By default, during a swap operation, the IEEE 802.1p bits of the VLAN tag remain unchanged. When the **swap-by-poppush** operation is enabled on a logical interface, the swap operation is treated as a **pop** operation followed by **push** operation. The **pop** operation removes the existing tag and the associated IEEE 802.1p bits and the push operation copies the inner VLAN IEEE 802.1p bits to the IEEE bits of the VLAN or VLANs being pushed. As a result, the IEEE 802.1p bits are inherited from the incoming transparent tag.

In effect, **swap-by-poppush** serves as a VLAN operation property and is used along with a **swap** or **swap-push** VLAN rewrite operation, indicating the nature of the swap operation being performed.

Related Documentation

- swap-by-poppush on page 1379
- *transparent*
- *Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag*
- Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag on page 588
- Understanding Transparent Tag Operations and IEEE 802.1p Inheritance on page 585

Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag

To classify incoming packets based on the IEEE 802.1p bits from the transparent tag, include the **transparent** statement at the **[edit class-of-service interfaces** *interface-name* **unit** *logical-unit-number* **classifiers ieee-802.1 vlan-tag]** hierarchy level.

Tagged Interface Example

The following example configuration specifies the classification based on the transparent VLAN tag.

```
edit
class-of-service {
  interfaces {
    ge-3/0/1 {
      unit 0 {
        classifiers {
          ieee-802.1 default vlan-tag transparent;
        }
      }
    }
  }
}
```

To configure Junos OS to inherit the IEEE 802.1p bits from the transparent tag, include the **swap-by-poppush** statement at the **[edit interfaces** ***interface-name*** **unit** ***logical-unit-number*]** hierarchy level.

The following is a configuration to swap and push VLAN tags and allow inheritance of the IEEE 802.1p value from the transparent VLAN tag in incoming packets.

```
edit
  ge-3/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 100;
    swap-by-poppush;
    input-vlan-map {
      swap-push;
      tag-protocol-id 0x9100;
      inner-tag-protocol-id 0x9100;
      vlan-id 500;
      inner-vlan-id 400;
    }
```



```
        output-vlan-map {
            pop-swap;
            inner-vlan-id 100;
            inner-tag-protocol-id 0x88a8;
        }
    }
}
```

The **swap-by-poppush** statement causes a swap operation to be done as a pop followed by a push operation. So for the outer tag, the incoming S-Tag is popped and a new tag is pushed. As a result, the S-Tag inherits the IEEE 802.1p bits from the transparent tag. The inner tag is then pushed, which results in the inner tag inheriting the IEEE 802.1p bits from the transparent tag.

Untagged Interface Example

The following is a configuration to push two VLAN tags and allow inheritance of the IEEE 802.1p value from the transparent VLAN tag in the incoming packet.

```
[edit]
ge-3/0/1 {
    encapsulation ccc;
    unit 0 {
        input-vlan-map {
            push-push;
            tag-protocol-id 0x9100;
            inner-tag-protocol-id 0x9100;
            vlan-id 500;
            inner-vlan-id 400;
        }
        output-vlan-map{
            pop-pop;
        }
    }
}
```

No additional configuration is required to inherit the IEEE 802.1p value, as the **push** operation inherits the IEEE 802.1p values by default.

The following configuration specifies the classification based on the transparent VLAN tag.

```
[edit]
class-of-service {
    interfaces {
        ge-3/0/1 {
            unit 0 {
                classifiers {
                    ieee-802.1 default vlan-tag transparent;
                }
            }
        }
    }
}
```

Related Documentation

- *transparent*
- swap-by-poppush on page 1379

- *Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag*
- Understanding swap-by-poppush on page 587
- Understanding Transparent Tag Operations and IEEE 802.1p Inheritance on page 585

PART 3

Operation, Administration, and Management (OAM) for Ethernet Interfaces

CHAPTER 31

Configuring IEEE 802.1ag OAM Connectivity-Fault Management

Ethernet Operations, Administration, and Maintenance

This topic provides an overview to help you effectively configure Ethernet Operations, Administration, and Maintenance (OAM) on a network of Juniper Networks® MX Series 3D Universal Edge Routers. For more information about configuring OAM parameters on Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

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

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

OAM can provide simple link-level information, provide performance statistics, or track end-to-end connectivity across the network. Simple link fault management (LFM) for Ethernet links is defined in IEEE 802.3ah.

IEEE 802.1ag OAM is supported on untagged, single tagged, and stacked VLAN interfaces.

Ethernet OAM functions are implemented as:

- Fault detection and notification (provided by continuity check messages)
- Path discovery (provided by the linktrace protocol)
- Fault isolation, verification, and recovery (isolation and verification are provided by a combination of protocols, while recovery is the function of protocols such as spanning tree)

The loopback protocol used in Ethernet OAM is modeled on the standard IP ping. After a fault is detected, the loopback protocol performs fault verification and isolation under the direction of a network operator.

The loopback is performed using request and response message pairs. A unicast loopback message is generated by a maintenance endpoint (MEP), and a loopback reply is generated by the destination maintenance intermediate point (MIP) or MEP.

The target MAC address is learned by the continuity check protocol or linktrace protocol. The loopback message's packet is always forwarded to a unique port by the originating MEP, as determined by a MAC table lookup or the MEP interface MAC address.

The target MIP or MEP generates a unicast loopback reply in response to the received loopback message. The loopback message follows the same path as a data packet, and intermediate bridges simply forward the packet to the destination MIP or MEP.

Related Documentation

- Ethernet OAM Connectivity Fault Management
- Example: Configuring Ethernet CFM on Bridge Connections
- Example: Configuring Ethernet CFM on Physical Interfaces

Ethernet OAM Connectivity Fault Management

The most complete connectivity fault management (CFM) is defined in IEEE 802.1ag. This topic emphasizes the use of CFM in a Metro Ethernet environment.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol. Similar to IP traceroute, this protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.
- Fault isolation using the loopback protocol. Similar to IP ping, this protocol works with the continuity check protocol during troubleshooting.

CFM partitions the service network into various administrative domains. For example, operators, providers, and customers might be part of different administrative domains.

Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible. Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outermost domains are assigned a higher level than the innermost domains.

Customer end points have the highest maintenance domain level. In a CFM maintenance domain, each service instance is called a maintenance association. A *maintenance association* can be thought as a full mesh of maintenance endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages.

There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

MEPs can be *up MEPs* or *down MEPs*. A link can connect a MEP at level 5 to a MEP at level 7. The interface at level 5 is an up MEP (because the other end of the link is at MEP level 7), and the interface at level 7 is a down MEP (because the other end of the link is at MEP level 5).

In a Metro Ethernet network, CFM is commonly used at two levels:

- By the service provider to check the connectivity among its provider edge (PE) routers
- By the customer to check the connectivity among its customer edge (CE) routers

NOTE: **The configured customer CFM level must be greater than service provider CFM level.**

In many Metro Ethernet networks, CFM is used to monitor connectivity over a VPLS and bridge network.

NOTE: **In ACX Series routers, VPLS is supported only on ACX5048 and ACX5096 routers.**

Related Documentation

- Ethernet Operations, Administration, and Maintenance on page 594
- Example: Configuring Ethernet CFM on Bridge Connections
- Example: Configuring Ethernet CFM on Physical Interfaces

IEEE 802.1ag OAM Connectivity Fault Management Overview

Ethernet interfaces on M7i and M10i routers with the Enhanced CFEB (CFEB-E) and on M120, M320, MX Series, T Series, and PTX Series routers support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). The goal of CFM is to monitor an Ethernet network that may comprise one or more service instances. Junos OS supports IEEE 802.1ag connectivity fault management.

NOTE: **MX Series Virtual Chassis does not support distributed inline connectivity fault management.**

ACX Series routers support CFM on aggregated Ethernet interfaces with continuity check interval of 100 milliseconds or higher.

In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces. Connectivity fault management (CFM) sessions operate in distributed mode and are processed on the Flexible PIC Concentrator (FPC) on aggregated Ethernet interfaces. As a result, graceful Routing Engine switchover (GRES) is supported on aggregated Ethernet interfaces. In releases before Junos OS Release 13.3, CFM sessions operate in centralized mode and are processed on the Routing Engine. However, CFM sessions are not supported on aggregated Ethernet interfaces if the interfaces that form the aggregated Ethernet bundle are in mixed mode. CFM sessions with a continuity check message (CCM) interval of 10 milliseconds are not supported over aggregated Ethernet interfaces.

CFM sessions are distributed by default. All CFM sessions must operate in either only distributed or only centralized mode. A mixed operation of distributed and centralized modes for CFM sessions is not supported. To disable the distribution of CFM sessions on aggregated Ethernet interfaces and make the sessions operate in centralized mode, include the **no-aggregate-delegate-processing** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

NOTE: As a requirement for Ethernet OAM 802.1ag to work, distributed periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine. You can only disable PPM on the Packet Forwarding Engine. To disable PPM on the PFE, include the **ppm no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

NOTE:

- CFM sessions are supported on aggregated Ethernet interfaces if the interfaces that form the aggregated Ethernet bundle are in mixed mode when the **no-aggregate-delegate-processing** command is enabled.
- Starting in Junos OS Release 14.2, , for CFM sessions in centralized mode, we recommend that you configure a maximum of 40 CFM sessions with continuity check message (CCM) interval of 100 milliseconds (100 ms) or a maximum of 400 CFM sessions with CCM interval of 1 second (1 s). If CFM sessions are configured beyond this limit, CFM might not work as expected. You might observe issues when the state of multiple links change or when the line cards are restarted.

  Note that these limits have been derived by considering a protocol data unit (PDU) load of 400 packets per second (pps) on the Routing Engine. This limit varies depending on the Routing Engine load. If the Routing Engine experiences heavy load, expect some variations to this limit.

Starting in Junos OS Release 10.3, on interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces.Starting in Junos OS Release 12.3, CFM does not support Multichassis Link Aggregation (MC-LAG). Do not configure the **mc-ae** statement when you configure CFM.

Starting in Junos OS Release 11.3, on T Series and M320 routers, CFM is not supported on interfaces configured with CCC encapsulation. If you configure CFM, the system displays the following message: “**MEPs cannot be configured on ccc interface on this platform**”.

Network entities such as operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management, perform end-to-end monitoring, and still avoid security breaches.

Starting in Junos OS Release 17.4, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces. Pseudowire service interfaces support configuring of subscriber interfaces over MPLS pseudowire termination. Termination of subscriber interfaces over PW enables network operators to extend their MPLS domain from the Access/Aggregation network to the service edge and use uniform MPLS label provisioning for a larger portion of their network.

NOTE: The CFM MIP session is supported only on the pseudowire services interface and not on the pseudowire services tunnel interface.

IEEE 802.1ag OAM supports graceful Routing Engine switchover (GRES). IEEE 802.1ag OAM is supported on untagged, single tagged, and stacked VLAN interfaces.

- Connectivity Fault Management Key Elements on page 598
- Best Practices for Configuring 802.1ag Ethernet OAM for VPLS on page 599

Connectivity Fault Management Key Elements

Figure 37 on page 598 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 37: Relationship Among MEPs, MIPs, and Maintenance Domain Levels

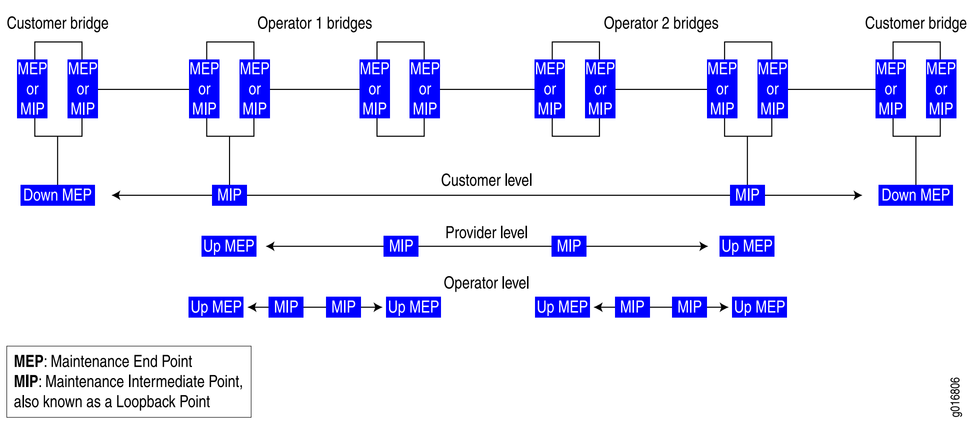

NOTE: On ACX Series routers, the maintenance intermediate points (MIP) are supported only on the ACX5048 and ACX5096 routers.

A maintenance association is a set of MEPs configured with the same maintenance association identifier and maintenance domain level. Figure 38 on page 599 shows the hierarchical relationships between the Ethernet bridge, maintenance domains, maintenance associations, and MEPs.

Figure 38: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs

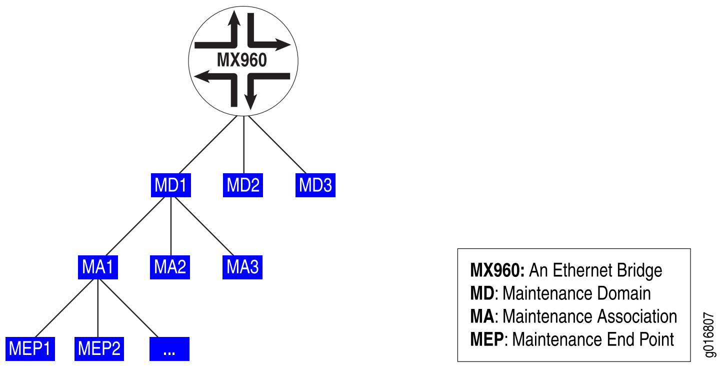

Best Practices for Configuring 802.1ag Ethernet OAM for VPLS

BEST PRACTICE: The logical interfaces in a VPLS routing instance may have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. VLAN normalization is effectively VLAN translation wherein the VLAN tags of the received packet need to be translated if they are different than the normalized VLAN tags. Configuration is described starting in "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 596 and you should further observe the additional requirements described in this section.

For MX Series routers, the normalized VLAN is specified using one of the following configuration statements in the VPLS routing instance:

- vlan-id *vlan-number*
- vlan-id none
- vlan-tags outer *outer-vlan-number* inner *inner-vlan-number*

You must configure vlan-maps explicitly on all interfaces belonging to the routing instance.

The following forwarding path considerations must be observed:

- Packet receive path:

- This is the forwarding path for packets received on the interfaces.
- 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.

- Packet transmit path:
  - The JUNOS Software uses the router's hardware-based forwarding for CPU-generated packets.
  - For Down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - In MX series routers, for Up MEPs, the packet must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packet to be forwarded with this flood route.
  - The router also uses implicit-based forwarding for CPU generated packets. The result is for the flood next hop tied to the flood route to be tied to the filter term. The filter term uses match criteria to correctly identify the host- generated packets.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces.
14.2	Starting in Junos OS Release 14.2, , for CFM sessions in centralized mode, we recommend that you configure a maximum of 40 CFM sessions with continuity check message (CCM) interval of 100 milliseconds (100 ms) or a maximum of 400 CFM sessions with CCM interval of 1 second (1 s).
12.3	Starting in Junos OS Release 12.3, CFM does not support Multichassis Link Aggregation (MC-LAG). Do not configure the **mc-ae** statement when you configure CFM.
11.3	Starting in Junos OS Release 11.3, on T Series and M320 routers, CFM is not supported on interfaces configured with CCC encapsulation.
10.3	Starting in Junos OS Release 10.3, on interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces.
9.3	In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces.

Related Documentation

- connectivity-fault-management on page 1113

Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36

Junos OS release 16.1R1 and later supports performance monitoring that is compliant with Technical Specification MEF 36. Technical Specification MEF 36 specifies the performance monitoring MIB. The performance monitoring MIB is required to manage service operations, administration, and maintenance (OAM) implementations that satisfy the Service OAM requirements and framework specified in MEF 17 and MEF 35, the management objects specified in MEF 7.1, and the performance monitoring functions defined in ITU-T Y.1731 and IEEE 802.1ag.

You can enable MEF-36-compliant performance monitoring by configuring the measurement-interval statement at the **[edit protocols oam ethernet cfm performance-monitoring]** hierarchy level.

When MEF-36-compliant performance monitoring is enabled:

- An SNMP get next request for a variable might not fetch the current value unless an SNMP walk is performed before performing the get next request. This limitation applies only to the current statistics for delay measurement, loss measurement, and synthetic loss measurement.
- The output for the field **Current delay measurement statistics** might display a measurement interval of 0 (zero) and an incorrect timestamp until the first cycle time has expired.
- Supported data TLV size for performance monitoring protocol data units (PDUs) is 1386 bytes when MEF-36-compliant performance monitoring is enabled. The TLV size is 1400 bytes in legacy mode.
- The maximum configurable value for the lower threshold bin is 4,294,967,294.

- Frame loss ratio (FLR) is excluded in loss measurements during period of unavailability for synthetic loss measurement only. In case of loss measurement, FLR is included even during period of unavailability.
- During a period of loss of continuity (adjacency down), although SOAM PDUs are not sent, FLR and availability calculations are not stopped. These calculations are performed with the assumption of 100% loss.
- The number of SOAM PDUs that are sent during the first measurement interval might be less than expected. This is because of a delay in detecting the adjacency state at the performance monitoring session level.
- The number of SOAM PDUs transmitted during a measurement interval for a cycle time of 100 ms might not be accurate. For example, in a measurement interval of two minutes with a cycle time 100 ms, the SOAM PDUs transmitted might be in the range of 1198—2000.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- **measurement-interval on page 1032**

Junos OS Support for Chassis ID TLV

In Release 16.1R2 and later, you can configure Junos OS to send the sender ID TLV along with the packets. The sender ID TLV is an optional TLV that is sent in continuity check messages (CCMs), loopback messages, and Link Trace Messages (LTMs), as specified in the IEEE 802.1ag standard. The sender ID TLV contains the chassis ID, which is the unique, CFM-based MAC address of the device, and the management IP address, which is an IPv4 or an IPv6 address.

The value of the **length** field in the TLV indicates whether or not the TLV contains the chassis ID information. The possible values for the **length** field are zero (**0**) or any valid number, which indicates the absence or presence of chassis ID information in the TLV, respectively.

You can enable Junos OS to send the sender ID TLV at the global level by using the **set protocols oam ethernet connectivity-fault-management sendid-tlv send-chassis-tlv** command. If the sender ID TLV is configured at the global level, then the default maintenance domain, maintenance association, and the maintenance association intermediate point (MIP) half function inherit this configuration.

You can also configure the sender ID TLV at the following hierarchy levels:

- **[edit protocols oam ethernet connectivity-fault-management]**.
- **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name* maintenance-association *maintenance-association-name* continuity-check]**.

The sender ID TLV configuration at the maintenance-association level takes precedence over the global-level configuration.

NOTE: The sender ID TLV is supported only for 802.1ag PDUs and is not supported for performance monitoring protocol data units (PDUs).

Release History Table

Release	Description
16.1	In Release 16.1R2 and later, you can configure Junos OS to send the sender ID TLV along with the packets.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

Creating a Maintenance Domain

To enable connectivity fault management (CFM) on an Ethernet interface, you must first configure a maintenance domain and specify the name of the maintenance domain. You can also specify the format of the name. For instance, if you specify the name format to be domain name service (DNS) format, you can specify the name of the maintenance domain as www.juniper.net. The default name format is ASCII character string.

NOTE: For logical interfaces, the maintenance domain name must be unique across logical systems. If you configure the same maintenance domain name across logical systems, then you receive the following error message: **error: configuration check-out failed.**

During the creation of the maintenance domain, you can also specify the maintenance domain level. The maintenance domain level indicates the nesting relationship between various maintenance domains. The maintenance domain level is embedded in each of the CFM frames.

To create a maintenance domain:

1. In configuration mode, create a maintenance domain by specifying the name and the name format at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

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

   NOTE: If you configure the maintenance domain name length greater than 45 octet, then the following error message is displayed: **error: configuration check-out failed.**

2. Specify the maintenance domain level by specifying the value at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@host# set maintenace-domain md-name level number
   ```

Related Documentation

- **connectivity-fault-management on page 1113**
- **maintenance-domain on page 1030**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- **name-format on page 1036**
- **level on page 1024**

Configuring Maintenance Intermediate Points (MIPs)

MX Series routers support maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag CFM protocol at a bridge-domain level. This enables you to define a maintenance domain for each default level. The MIPs names are created as **default-*level-number*** at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level. Use the **bridge-domain**, **instance**, **virtual-switch**, and **mip-half-function** MIP options to specify the MIP configuration.

Use the **show oam ethernet connectivity-fault-management mip (bridge-domain | instance-name | interface-name)** command to display the MIP configurations.

To configure the maintenance intermediate point (MIP):

1. Configure a bridge domain under a user-defined virtual switch by specifying the **virtual-switch** statement and the name of the user-defined virtual switch, at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* default-*x*]** hierarchy level.

NOTE: A bridge domain must be specified by name only if it is configured by including the **vlan-id** statement under the **virtual-switch** statement. If a bridge domain is configured with a range of VLAN IDs, then the VLAN IDs must be explicitly listed after the bridge domain name.

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* default-*x*]
user@host# set virtual-switch *virtual-switch-name* bridge-domain *bridge-domain-name* vlan-id *value*

NOTE: You can also configure the bridge domain for the default virtual switch by including the **bridge-domain** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

2. Configure the VPLS routing instance for the default maintenance domain.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]
   user@host# set instance *instance-name*

3. Configure the maintenance intermediate point (MIP) half function to divide the MIP fuctionality into two unidirectional segments to improve network coverage by increasing the number of MIPs that are monitored. The MIP half function also responds to loop-back and link-trace messages to identify faults.

   NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations be the same.

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

Related Documentation

- bridge-domain on page 1092
- connectivity-fault-management on page 1113
- instance on page 1017
- mip-half-function on page 1035
- virtual-switch on page 1431
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603

Configuring Maintenance Association Intermediate Points in ACX Series

Maintenance Intermediate Point (MIP) provides monitoring capability of intermediate points for services such as Layer 2 bridging, Layer 2 circuit, and Layer 2 VPN. ACX5048 and ACX5096 routers support MIPs for the Ethernet OAM 802.1ag CFM protocol. Use the bridge-domain, interface, and mip-half-function MIP options to specify the MIP configuration.

NOTE: ACX5048 and ACX5096 routers do not support MIP configuration on VPLS services.

NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations be the same.

To display MIP configurations, use the **show oam ethernet connectivity-fault-management mip (bridge-domain | instance-name | interface-name)** command.

The following MIP configurations are supported in ACX5048 and ACX5096 routers:

- MIP with with bridge domain
- MIP with circuit cross-connect (CCC)
- MIP with bridge domain when maintenance association end point is configured
- MIP with CCC when maintenance association end point is configured

The following sections describe MIP configuration:

Configuring the Maintenance Domain Bridge Domain

To configure the bridge domain, include the **vlans** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***maintenance-domain-name*****]** hierarchy level.

NOTE: The Layer 2 CLI configurations and show commands for ACX5048 and ACX5096 routers differ compared to other ACX Series routers. For more information, see *Understanding Layer 2 Next Generation Mode on ACX Series Routers*.

Configuring the Maintenance Domain MIP Half Function

MIP Half Function (MHF) divides MIP functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loopback and linktrace messages to help isolate faults.

Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the *MIP half function* value for all maintenance domains and maintenance associations be the same. To configure the MIP half function, include the mip-half-function statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***maintenance-domain-name*****]** hierarchy level.

Configuring the Maintenance Association Intermediate Points with Bridge Domain

In ACX5048 and ACX5096 routers, you can configure the MIP with bridge domain. The following is a sample to configure the MIP with bridge domain:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain default-6 {
        vlan bd1;
```

```
                mip-half-function default;
            }
        }
    }
}
```

Configuring the Maintenance Association Intermediate Points with Circuit Cross-Connect

In ACX5048 and ACX5096 routers, you can configure the MIP with circuit cross-connect (CCC). The following is a sample to configure the MIP with CCC:

```
[edit protocols]
oam {
    ethernet {
        connectivity-fault-management {
            maintenance-domain default-6 {
                interface xe-0/0/42.0;
                mip-half-function default;
            }
        }
    }
}
```

Configuring the Maintenance Association Intermediate Points with Bridge Domain when Maintenance Association End Point is Configured

In ACX5048 and ACX5096 routers, you can configure the MIP with bridge domain when a maintenance association end point (MEP) is configured. The following is a sample to configure the MIP with bridge domain when MEP is configured:

```
[edit protocols]
oam {
    ethernet {
        connectivity-fault-management {
            maintenance-domain md2 {
                level 5;
                mip-half-function default;
                maintenance-association ma2 {
                    continuity-check {
                        interval 1s;
                    }
                    mep 222 {
                        interface xe-0/0/42.0;
                        direction up;
                    }
                }
            }
        }
    }
}
```


Configuring the Maintenance Intermediate Points with Circuit Cross-Connect when Maintenance Association End Point is Configured

In ACX5048 and ACX5096 routers, you can configure the MIP with circuit cross-connect (CCC) when a maintenance association end point (MEP) is configured. The following is a sample to configure the MIP with CCC when MEP is configured:

```
[edit protocols]
oam {
   ethernet {
      connectivity-fault-management {
         maintenance-domain md2 {
            level 5;
            mip-half-function default;
            maintenance-association ma2 {
               continuity-check {
                  interval 1s;
               }
               mep 222 {
                  interface xe-0/0/42.0;
                  direction up;
               }
            }
         }
      }
   }
}
```

Related Documentation

- **bridge-domain on page 1092**
- **connectivity-fault-management on page 1113**
- **instance on page 1017**
- **mip-half-function on page 1035**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Creating a Maintenance Association on page 610
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- *Ethernet Interfaces Feature Guide for Routing Devices*

Creating a Maintenance Association

To create a maintenance association, include the **maintenance-association *ma-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

Maintenance association names can be in one of the following formats:

- As a plain ASCII character string
- As the VLAN identifier of the VLAN you primarily associate with the maintenance association
- As a two-octet identifier in the range from 0 through 65,535
- As a name in the format specified by RFC 2685

The default short name format is an ASCII character string.

To configure the maintenance association short name format, include the **short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id)** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level.

Related Documentation

- **connectivity-fault-management on page 1113**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655

Continuity Check Protocol Parameters Overview

The continuity check protocol is used for fault detection by maintenance end points (MEPs) within a maintenance association. The MEP periodically sends continuity check multicast messages. The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32.

The following list describes the continuity check protocol parameters you can configure:

- **interval**—Frequency of the continuity check messages (CCM) i.e time between the transmission of the CCM messages. You can specify 10 minutes (**10m**), 1 minute (**1m**), 10 seconds (**10s**), 1 second (**1s**), 100 milliseconds (**100ms**), or 10 milliseconds (**10ms**). The default value is 1 minute. For instance, if you specify the interval as 1 minute, the MEP sends the continuity check messages every minute to the receiving MEP.

  NOTE: For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can only disable PPM on the Packet Forwarding Engine. To disable PPM on the Packet Forwarding Engine, use the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

  Continuity check interval of 10 milliseconds is not supported for CFM sessions over a label-switched interface (LSI).

- **hold-interval**—Frequency at which the MEP database can be flushed, if no updates occur. Receiving MEPs use the continuity check messages to build a MEP database of all MEPs in the maintenance association. The frequency is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

  NOTE: Hold timer based flushing is applicable only for autodiscovered remote MEPs and not for statically configured remote MEPs.

  The hold interval logic runs a polling timer per CFM session level (not per remote MEP level) where the polling timer duration is equal to the configured hold time. When the polling timer expires, it deletes all the autodiscovered remote MEP entries which have been in the failed state for a time period equal to or greater than the configured hold time. If the remote MEP completes the hold time duration in the failed state, then flushing will not occur until the next polling timer expires. Hence remote MEP flushing may not happen exactly at the configured hold time.

- **loss-threshold**—Number of continuity check messages that can be lost before the router marks the MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- Managing Continuity Measurement Statistics on page 803

Configuring Continuity Check Protocol Parameters for Fault Detection

The continuity check protocol is used for fault detection by a maintenance association end point (MEP) within a maintenance association. A MEP periodically generates and responds to continuity check multicast messages. The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32. The receiving MEPs use the continuity check messages (CCMs) to build a MEP database of all MEPs in the maintenance association.

To configure continuity check protocol parameters:

1. Specify the time to wait in minutes before flushing the MEP database, if no updates occur, with a value from 1 minute through 30,240 minutes. The default value is 10 minutes.

   NOTE: **Flushing based on the hold timer is applicable only for autodiscovered remote MEPs and not for statically configured remote MEPs.**

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check]
   user@host# set hold-interval *minutes*

2. Specify the time to wait (duration) between the transmissions of CCMs. The duration can be one of the following values: 10 minutes (10m), 1 minute (1m), 10 seconds (10s),

1 second (1s), 100 milliseconds (100ms), or 10 milliseconds (10ms). The default value is 1 minute.

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check]
user@host# set interval *duration*

3. Specify the number of continuity check messages that can be lost before the router marks the MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

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

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Continuity Check Protocol Parameters Overview on page 611
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- Managing Continuity Measurement Statistics on page 803
- Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades on page 800

Configuring a MEP to Generate and Respond to CFM Protocol Messages

A maintenance association end point (MEP) refers to the boundary of a domain. A MEP generates and responds to connectivity fault management (CFM) protocol messages. You can configure multiple up MEPs for a single combination of maintenance association ID and maintenance domain ID for interfaces belonging to a particular VPLS service or a bridge domain. You can configure multiple down MEPs for a single instance of maintenance domain identifier and maintenance association name to monitor services provided by Virtual Private LAN service (VPLS), bridge domain, circuit cross-connect (CCC), or IPv4 domains.

For layer 2 VPNs routing instances (local switching) and EVPN routing instances, you can also configure multiple up MEPs for a single combination of maintenance association ID and maintenance domain ID on logical interfaces.. The logical interface can be configured on different devices or on the same device. To support multiple up MEPs on two IFLs, enhanced IP network services must be configured for the chassis.

You can enable automatic discovery of a MEP. With automatic discovery a MEP is enabled to accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association. if automatic discovery is not enabled, the remote MEPs must be configured. If the remote MEP is not configured, the CCMs from the remote MEP are treated as errors.

Continuity measurement is provided by an existing continuity check protocol. The continuity for every remote MEP is measured as the percentage of time that remote MEP was operationally up over the total administratively enabled time. Here, the operational uptime is the total time during which the CCM adjacency is active for a particular remote MEP and the administrative enabled time is the total time during which the local MEP is active. You can also restart the continuity measurement by clearing the currently measured operational uptime and the administrative enabled time.

- Configuring a Maintenance Association End Point (MEP) on page 614
- Configuring a remote Maintenance Association End Point (MEP) on page 616

Configuring a Maintenance Association End Point (MEP)

To configure a maintenance association end point:

1. Specify an ID for the MEP at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name***]. You can specify any value from 1 through 8191.

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

2. Enable maintenance end point automatic discovery so the MEP can accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association.

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

3. Specify the direction in which the CCM packets are transmitted for the MEP. You can specify up or down. If you specify the direction as up, CCMs are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on the MEP. If you specify the direction as down, CCMs are transmitted only out of the interface configured on the MEP.

NOTE: Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.

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

NOTE: Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all other routers and switches, you must continue to configure the **no-control-word** statement at the **[edit routing-instances *routing-instance-name* protocols l2vpn]** or **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

4. Specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface. On MX Series routers, the MEP can be attached to a specific VLAN of a trunk interface.

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

5. Specify the IEEE 802.1 priority bits that are used by continuity check and link trace messages. You can specify a value from through 7 as the priority.

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

6. Specify the lowest priority defect that generates a fault alarm whenever CFM detects a defect. Possible values include: all -defects, err-xcon, mac-rem-err-xcon, no-defect, rem-err-xcon, and xcon.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  domain-name maintenance-association ma-name mep mep-id]
user@host# set lowest-priority-defect mac-rem-err-xcon
```

7. Specify the ID of the remote MEP at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name***

maintenance-association *ma-name* mep *mep-id*]. You can specify any value from 1 through 8191.

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

See Also

- auto-discovery on page 994
- direction on page 1010
- lowest-priority-defect on page 1028
- priority on page 1044

Configuring a remote Maintenance Association End Point (MEP)

To configure a remote maintenance association end point:

1. Configure the remote MEP by specifying the MEP ID at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id***]. You can specify any value from 1 through 8191.

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

2. Specify the name of the action profile to be used for the remote MEP by including the **action-profile *profile-name*** statement at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id***]. The profile must be defined at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

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

3. Configure the remote MEP to detect initial loss of connectivity. By default, the MEP does not generate loss-of-continuity (LOC) defect messages. When you configure the **detect-loc** statement, a loss-of-continuity (LOC) defect is detected if no continuity check message is received from the remote MEP within a period equal to 3.5 times the continuity check interval configured for the maintenance association. If a LOC defect is detected, a syslog error message is generated.

NOTE: When you configure connectivity-fault management (CFM) along with **detect-loc**, any **action-profile** configured to bring down the interface is executed if continuity check message is not received . However, the **action-profile** is not executed if you have not configured **detect-loc** and continuity check message is not received.

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*mep *mep-id* remote-mep *remote-mep-id*]
user@host# set detect-loc

See Also

- action-profile on page 993
- detect-loc on page 1009
- remote-mep on page 1048

Release History Table

Release	Description
12.3	Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs.

Related Documentation

- connectivity-fault-management on page 1113
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- Configuring Service Protection for VPWS over MPLS Using the MEP Interface on page 620

Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging Service Level Agreements (SLAs). By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

Before you can perform Ethernet frame delay measurements on MX Series routers, you must have done the following:

- Configured Ethernet OAM and CFM correctly
- Prepared the measurement between two compatibly configured MX Series routers
- Enabled the distributed periodic packet management daemon (ppmd)
- Avoided trying to perform Ethernet frame delay measurement on aggregated Ethernet or pseudowire interfaces, which are not supported
- Made sure the hardware-assisted timestamping is supported if that feature is configured

At the end of this configuration, you create two MX Series routers that can perform and display Ethernet frame delay measurements on Ethernet interfaces using optional hardware timestamping. By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

To configure hardware-assisted timestamping:

1. To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

   ```
   [edit]
   protocols {
     oam {
       ethernet {
         connectivity-fault-management {
           performance-monitoring {
             hardware-assisted-timestamping; # Enable timestamping in hardware.
           }
         }
       }
     }
   ```

2. Ethernet frame delay measurement requires that distributed PPMD is enabled. Before you can gather statistics for Ethernet frame delay measurement, you must make sure that PPMD is configured properly. Without distributed PPMD, delay measurement results are not valid.

To perform Ethernet frame delay measurement, make sure that the following configuration statement is *NOT* present:

```
[edit routing-options]
ppm {
   no-delegate-processing; # This turns distributed PPMD OFF.
}
```

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Triggering an ETH-DM Session on page 852
- Viewing ETH-DM Statistics on page 853
- Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 817
- Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 822
- Configuring ETH-DM with Untagged Interfaces on page 826

Configuring Service Protection for VPWS over MPLS Using the MEP Interface

You can enable service protection for a virtual private wire service (VPWS) over MPLS by specifying a working path or protect path on the MEP. Service protection provides end-to-end connection protection of the working path in the event of a failure.

To configure service protection, you must create two separate transport paths—a working path and a protect path. You can specify the working path and protect path by creating two maintenance associations. To associate the maintenance association with a path, you must configure the **interface** statement for the MEP within the maintenance association and specify the path as working or protect.

NOTE: **If the path is not specified, the session monitors the active path.**

Table 63 on page 620 describes the available service protection options.

Table 63: Service Protection Options

Option	Description
working	Specifies the working path.
protect	Specifies the protect path.

In this configuration, we enable service protection for the VPWS service. The CCM session is configured for the working path and references the CCM session configured for the protect path using the **protect-maintenance-association** statement. The name of the protect transport path for the maintenance association is configured and associated with the maintenance association for the working path.

To configure service protection for VPWS over MPLS:

1. In configuration mode, create a maintenance domain by specifying the name and the name format at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

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

   NOTE: **If you configure the maintenance domain name length greater than 45 octet, then the following error message is displayed: error: configuration check-out failed.**

2. Specify the maintenance domain level by specifying the value at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

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


user@host# set maintenace-domain *md-name* level *number*

3. Create a maintenance association for the working path by specifying the name and the short name format at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name***] hierarchy level.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name*]
   user@host# set maintenance-association test-ma short-name-format *option*

4. Specify the maintenance association name used for connection protection and the name of the automatic-protection-switching profile (aps-profile) at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name***] hierarchy level.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]
   user@host# set protect-maintenance-association *ma-name* aps-profile *aps-profile-name*

5. Specify the time to wait between transmissions of continuity check messages at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check**] hierarchy level. The duration can be one of the following values: 10 minutes(10m), 1 minute(1m), 10 seconds(10s), 1 second(1s), 100 milliseconds(100ms), or 10 milliseconds(10ms). The default value is 1 minute.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check]
   user@host# set interval option

6. Specify an ID for the MEP at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name***]. You can specify any value from 1 through 8191.

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

7. Enable maintenance end point automatic discovery so the MEP can accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association.

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

8. Specify the direction in which the CCM packets are transmitted for the MEP. You can specify up or down. If you specify the direction as up, CCMs are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on the MEP. If you specify the direction as down, CCMs are transmitted only out of the interface configured on the MEP.

NOTE: Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.

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

NOTE: Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all other routers and switches, you must continue to configure the **no-control-word** statement at the **[edit routing-instances *routing-instance-name* protocols l2vpn]** or **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

9. Specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface. On MX Series routers, the MEP can be attached to a specific VLAN of a trunk interface. Also, specify the transport path as working.

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

10. Create a maintenance association for the protection path by specifying the name and the short name format at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name***] hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name]
user@host# set maintenance-association ma-name short-name-format option
```

11. Specify the time to wait between transmissions of continuity check messages at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check**] hierarchy level. The duration can be one of the following values: 10 minutes(10m), 1 minute(1m), 10 seconds(10s), 1 second(1s), 100 milliseconds(100ms), or 10 milliseconds(10ms). The default value is 1 minute.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name continuity-check]
```


user@host# set interval option

12. Specify an ID for the MEP at the [**edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name***]. You can specify any value from 1 through 8191.

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

13. Enable maintenance end point automatic discovery so the MEP can accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association.

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

14. Specify the direction in which the CCM packets are transmitted for the MEP. You can specify up or down. If you specify the direction as up, CCMs are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on the MEP. If you specify the direction as down, CCMs are transmitted only out of the interface configured on the MEP.

    NOTE: Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.

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

    NOTE: Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all other routers and switches, you must continue to configure the **no-control-word** statement at the [**edit routing-instances** ***routing-instance-name*** **protocols l2vpn**] or [**edit protocols l2circuit neighbor** ***neighbor-id*** **interface** ***interface-name***] hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

15. Specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface. On MX Series routers, the MEP can be attached to a specific VLAN of a trunk interface. Also, specify the transport path as working.

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

Release History Table

Release	Description
12.3	Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs.
12.3	Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs.

Related Documentation

- **auto-discovery on page 994**
- **interval on page 1020**
- **name-format on page 1036**
- **protect-maintenance-association on page 1046**
- **short-name-format on page 1050**
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

You can create a connectivity fault management (CFM) action profile to define event flags and thresholds to be monitored. You can also specify the action to be taken when any of the configured events occur. When the CFM events occur, the router performs the corresponding action based on your specification. You can configure one or more events in the action profile. Alternatively, you can configure an action profile and specify default actions when connectivity to a remote maintenance association endpoint (MEP) fails.

NOTE: You cannot configure multiple actions at this time. Only one action can be configured. This limitation affects both the **action** and **clear-action** statements.

To configure the CFM action profile:

1. In configuration mode, at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level, specify the name of the action profile and the CFM event(s). You can configure more than one event in the action profile. Possible events include: interface-status-tlv, port-status-tlv, adjacency-loss, RDI.

   [edit protocols oam ethernet connectivity-fault-management]
   user@host# set action-profile *profile-name* event [*event1*, *event2*, *event3*..]

2. Specify the action to be taken by the router when the event occurs. The action is triggered when the event occurs. If you have configured more than one event in the action profile, it is not necessary for all events to occur to trigger the action.

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

3. Specify the default action to be taken by the router when connectivity to a remote MEP fails. If no action is configured, no action is taken.

   NOTE: Associating an action profile with the **interface-down** action on an up MEP CFM session running over a circuit cross-connect (CCC) interface (l2circuit/l2vpn) is not advisable and can result in a deadlock situation.

   [edit protocols oam ethernet connectivity-fault-management]
   user@host# set action-profile *profile-name* default-actions *action*

Related Documentation

- event (CFM) on page 1012
- default-actions on page 1005
- connectivity-fault-management on page 1113
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

- Creating a Maintenance Domain on page 603
- Creating a Maintenance Association on page 610
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

Configuring Linktrace Protocol in CFM

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between an MEP and an MIP under the same maintenance domain. The linktrace protocol enables you to configure the time to wait for a response. If no response is received for a linktrace request message, the request and response entries are deleted after the interval expires. You can also configure the number of linktrace reply entries to be stored for the corresponding linktrace request.

The operation of IEEE 802.1ag linktrace request and response messages is similar to the operation of Layer 3 **traceroute** commands. For more information about the **traceroute** command, see the *Junos OS Administration Library*.

To configure the linktrace protocol:

1. Configure the time to wait for a linktrace response at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level. You can specify the value in minutes or seconds. The default value is 10 minutes.

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

2. Configure the number of linktrace reply entries to be stored per linktrace request. You can specify a value from 1 through 500. The default value is 100.

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

Related Documentation

- **age on page 1075**
- **path-database-size on page 1037**
- **connectivity-fault-management on page 1113**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625

Configuring Ethernet Local Management Interface

Ethernet Local Management Interface Overview

Gigabit Ethernet (**ge**), 10-Gigabit Ethernet (**xe**), and Aggregated Ethernet (**ae**) interfaces support the Ethernet Local Management Interface (E-LMI).

NOTE: On MX Series routers, E-LMI is supported on Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), and Aggregated Ethernet (ae) interfaces configured on MX Series routers with DPC only.

The E-LMI specification is available at the Metro Ethernet Forum. E-LMI procedures and protocols are used for enabling automatic configuration of the customer edge (CE) to support Metro Ethernet services. The E-LMI protocol also provides user-to-network interface (UNI) and Ethernet virtual connection (EVC) status information to the CE. The UNI and EVC information enables automatic configuration of CE operation based on the Metro Ethernet configuration.

The E-LMI protocol operates between the CE device and the provider edge (PE) device. It runs only on the PE-CE link and notifies the CE of connectivity status and configuration parameters of Ethernet services available on the CE port. The scope of the E-LMI protocol is shown in Figure 39 on page 627.

Figure 39: Scope of the E-LMI Protocol

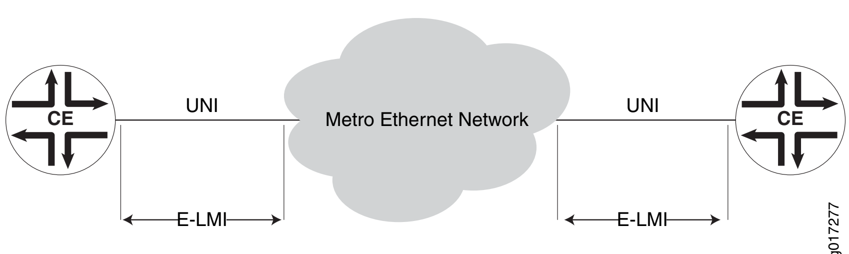

The E-LMI implementation on ACX and MX Series routers includes only the PE side of the E-LMI protocol.

E-LMI interoperates with an OAM protocol, such as Connectivity Fault Management (CFM), that runs within the provider network to collect OAM status. CFM runs at the provider maintenance level (UNI-N to UNI-N with up MEPs at the UNI). E-LMI relies on the CFM for end-to-end status of EVCs across CFM domains (SVLAN domain or VPLS).

The E-LMI protocol relays the following information:

- Notification to the CE of the addition/deletion of an EVC (active, not active, or partially active)
- Notification to the CE of the availability state of a configured EVC
- Communication of UNI and EVC attributes to the CE:
  - UNI attributes:
    - UNI identifier (a user-configured name for UNI)
    - CE-VLAN ID/EVC map type (all-to-one bundling, service multiplexing with bundling, or no bundling)
    - Bandwidth profile is not supported (including the following features):
      - CM (coupling mode)
      - CF (color flag)
      - CIR (committed Information rate)
      - CBR (committed burst size)
      - EIR (excess information rate)
      - EBS (excess burst size)
  - EVC attributes:
    - EVC reference ID
    - EVC status type (active, not active, or partially active)
    - EVC type (point-to-point or multipoint-to-multipoint)
    - EVC ID (a user-configured name for EVC)
    - Bandwidth profile (not supported)
  - CE-VLAN ID/EVC map

E-LMI on MX Series routers supports the following EVC types:

- Q-in-Q SVLAN (point-to-point or multipoint-to-multipoint)—Requires an end-to-end CFM session between UNI-Ns to monitor the EVS status.
- VPLS (BGP or LDP) (point-to-point or multipoint-to-multipoint)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.
- L2 circuit/L2VPN (point-to-point)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.

NOTE: **l2-circuit and l2vpn are not supported.**

The E-LMI protocol on ACX Series routers supports Layer 2 circuit and Layer 2 VPN EVC types and enables link-loss forwarding for pseudowire (Layer 2 circuit and Layer 2 VPN) services as follows:

- Interworking between the connectivity fault management (CFM) protocol and the E-LMI protocol for Layer 2 circuit and Layer 2 VPN.
  - End-to-end CFM session between UNIs to monitor EVC status.
  - In the case of pseudowire redundancy, CFM can be used to monitor active and backup pseudowire sessions. The EVC status is declared as down to CE devices only when both the active and backup pseudowire sessions go down.
- Interworking between remote defect indication (RDI) and E-LMI for Layer 2 circuit and Layer 2 VPN.
  - If a maintenance association end point (MEP) receives an RDI bit set in a continuity check message (CCM) frame, and if RDI fault detection is enabled in the EVC configuration at **[edit protocols oam ethernet evcs *evc-id* evc-protocol *cfm* management-domain *name* management-association *name* faults rdi]**, then the pseudowire is declared as down to CE routers through E-LMI.
- If an end-to-end CFM session does not exist between UNIs, the pseudowire (Layer 2 circuit or Layer 2 VPN) up and down state triggers an asynchronous EVC state change message to CE routers through E-LMI.

NOTE: ACX Series routers do not support E-LMI for Layer 2 services (bridging).

Configuring the Ethernet Local Management Interface

To configure E-LMI, perform the following steps:

Configuring an OAM Protocol (CFM)

For information on configuring the OAM protocol (CFM), see “IEEE 802.1ag OAM Connectivity Fault Management Overview” on page 596.

Assigning the OAM Protocol to an EVC

To configure an EVC, you must specify a name for the EVC using the **evcs*evc-id*** statement at the **[edit protocols oam ethernet]** hierarchy level. You can set the EVC protocol for monitoring EVC statistics to **cfm** or **vpls** using the **evc-protocol** statement and its options at the **[edit protocols oam ethernet evcs]** hierarchy level.

You can set the number of remote UNIs in the EVC using the **remote-uni-count** ***number*** statement at the **[edit protocols oam ethernet evcs evcs-protocol]** hierarchy level. The **remote-uni-count** defaults to 1. Configuring a value greater than 1 makes the EVC multipoint-to-multipoint. If you enter a value greater than the actual number of endpoints, the EVC status will display as partially active even if all endpoints are up. If you enter a **remote-uni-count** less than the actual number of endpoints, the status will display as active, even if all endpoints are not up.

You can configure an EVC by including the **evcs** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
evcs evc-id {
  evc-protocol (cfm (management-domain name management-association name ) | vpls
    (routing-instance name)) {
    remote-uni-count <number>; # Optional, defaults to 1
    multipoint-to-multipoint;
      # Optional, defaults to point-to-point if remote-uni-count is 1
  }
}
```

Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC

To configure E-LMI, include the **lmi** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
lmi {
  polling-verification-timer value;
    # Polling verification timer (T392), defaults to 15 seconds
  status-counter count; # Status counter (N393), defaults to 4
  interface name {
    evc evc-id {
      default-evc;
      vlan-list [ vlan-ids ];
    }
    evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
    polling-verification-time value; # Optional, defaults to global value
    status-counter count; # Optional, defaults to global value
    uni-id value; # Optional, defaults to interface-name
  }
}
```

You can set the status counter to count consecutive errors using the **status-counter** ***count*** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The status counter is used to determine if E-LMI is operational or not. The default value is 4.

You can set the **polling-verification-timer** ***value*** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The default value is 15 seconds.

You can enable an interface and set its options for use with E-LMI using the **interface** ***name*** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. Only **ge**, **xe**, and **ae** interfaces are supported. You can use the interface **uni-id** option to specify a name for the UNI. If **uni-id** is not configured, it defaults to the name variable of **interface** ***name***.

You can specify the CE-VLAN ID/EVC map type using the **evc-map-type type** interface option. The options are **all-to-one-bundling**, **bundling**, or **service-multiplexing**. Service multiplexing is with no bundling. The default type is **all-to-one-bundling**.

To specify the EVC that an interface uses, use the **evc *evc-id*** statement at the **[edit protocols oam ethernet lmi interface *name*]** hierarchy level. You can specify an interface as the default EVC interface using the **default-evc** statement at the **[edit protocols oam ethernet lmi interface *name* evc *evc-id*]** hierarchy level. All VIDs that are not mapped to any other EVCs are mapped to this EVC. Only one EVC can be configured as the default.

You can map a list of VLANs to an EVC using the **vlan-list *vlan-id-list*** statement at the **[edit protocols oam ethernet lmi interface *name* evc *evc-id*]** hierarchy level.

Example E-LMI Configuration

- Example Topology on page 631
- Configuring PE1 on page 631
- Configuring PE2 on page 633
- Configuring Two UNIs Sharing the Same EVC on page 634

Example Topology

Figure 40 on page 631 illustrates the E-LMI configuration for a point-to-point EVC (SVLAN) monitored by CFM. In this example, VLANs 1 through 2048 are mapped to **evc1** (SVLAN 100) and 2049 through 4096 are mapped to **evc2** (SVLAN 200). Two CFM sessions are created to monitor these EVCs.

Figure 40: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM

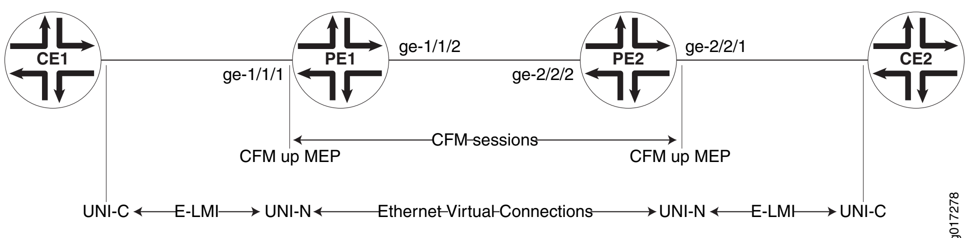

Configuring PE1

```
[edit]
interfaces {
   ge-1/1/1 {
      unit 0 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 1-2048;
         }
      }
      unit 1 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 2049-4096;
         }
```

```
        }
      }
      ge-1/1/2 {
        unit 0 {
          vlan-id 100;
          family bridge {
            interface-mode trunk;
            inner-vlan-id-list 1-2048;
          }
        }
        unit 1 {
          vlan-id 200;
          family bridge {
            interface-mode trunk;
            inner-vlan-id-list 2049-4096;
          }
        }
      }
    }
    protocols {
      oam {
        ethernet {
          connectivity-fault-management {
            maintenance-domain md {
              level 0;
              maintenance-association 1 {
                name-format vlan;
                mep 1 {
                  direction up;
                  interface ge-1/1/1.0 vlan 1;
                }
              }
              maintenance-association 2049 {
                name-format vlan;
                mep 1 {
                  direction up;
                  interface ge-1/1/1.1 vlan 2049;
                }
              }
            }
          }
          evcs {
            evc1 {
              evc-protocol cfm management-domain md management-association 1;
              remote-uni-count 1;
            }
            evc2 {
              evc-protocol cfm management-domain md management-association 2049;
              remote-uni-count 1;
            }
          }
          lmi {
            interface ge-1/1/1 {
              evc evc1 {
                vlan-list 1-2048;
              }
```



```
                    evc evc2 {
                        vlan-list 2049-4096;
                    }
                    evc-map-type bundling;
                    uni-id uni-ce1;
                }
            }
        }
    }
}
```

Configuring PE2

```
[edit]
interfaces {
    ge-2/2/1 {
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list 1-2048;
            }
        }
        unit 1 {
            family bridge {
                interface-mode trunk;
                vlan-id-list 2049-4096;
            }
        }
    }
    ge-2/2/2 {
        unit 0 {
            vlan-id 100;
            family bridge {
                interface-mode trunk;
                inner-vlan-id-list 1-2048;
            }
        }
        unit 1 {
            vlan-id 200;
            family bridge {
                interface-mode trunk;
                inner-vlan-id-list 2049-4095;
            }
        }
    }
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                maintenance-domain md {
                    level 0;
                    maintenance-association 1 {
                        name-format vlan;
                        mep 1 {
                            direction up;
```

```
                interface ge-2/2/1.0 vlan 1;
              }
            }
            maintenance-association 2049 {
              name-format vlan;
              mep 1 {
                direction up;
                interface ge-2/2/1.1 vlan 2049;
              }
            }
          }
        }
        evcs {
          evc1 {
            evc-protocol cfm management-domain md management-association 1;
            remote-uni-count 1;
          }
          evc2 {
            evc-protocol cfm management-domain md management-association 2049;
            uni-count 2;
          }
        }
        lmi {
          interface ge-2/2/1 {
            evc evc1 {
              vlan-list 1-2048;
            }
            evc evc2 {
              vlan-list 2049-4095;
            }
            evc-map-type bundling;
            uni-id uni-ce2;
          }
        }
      }
    }
  }
```

Configuring Two UNIs Sharing the Same EVC

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management { ...}
    evcs {
      evc1 {
        evc-protocol cfm management-domain md management-association 1;
        remote-uni-count 1;
      }
    }
    lmi {
      interface ge-2/2/1 {
        evc evc1 {
          vlan-list 0-4095;
        }
        evc-map-type all-to-one-bundling;
```



```
            uni-id uni-ce1;
        }
        interface ge-2/3/1 {
            evc evc1 {
                vlan-list 0-4095;
            }
            evc-map-type all-to-one-bundling;
            uni-id uni-ce2;
        }
      }
    }
  }
```

Related Documentation

- **connectivity-fault-management on page 1113**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655

Configuring Port Status TLV and Interface Status TLV

TLVs Overview

Type, Length, and Value (TLVs) are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and/or optional information in a PDU. TLVs are not aligned to any particular word or octet boundary. TLVs follow each other with no padding between them.

Table 64 on page 636 shows the TLV format and indicates if it is required or optional.

Table 64: Format of TLVs

Parameter	Octet (sequence)	Description
Type	1	Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.
Length	2–3	Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field.
Value	4	Length specified by the Length field. Optional. Not present if the Type field is 0 or if the Length field is 0.

Various TLVs for CFM PDUs

Table 65 on page 636 shows a set of TLVs defined by IEEE 802.1ag for various CFM PDU types. Each TLV can be identified by the unique value assigned to its type field. Some type field values are reserved.

Table 65: Type Field Values for Various TLVs for CFM PDUs

TLV or Organization	Type Field
End TLV	0
Sender ID TLV	1
Port Status TLV	2
Data TLV	3
Interface Status TLV	4
Reply Ingress TLV	5
Reply Egress TLV	6
LTM Egress Identifier TLV	7
LTR Egress Identifier TLV	8
Reserved for IEEE 802.1	9 to 30
Organization-Specific TLV	31
Defined by ITU-T Y.1731	32 to 63
Reserved for IEEE 802.1	64 to 255

Not every TLV is applicable for all types of CFM PDUs.

- TLVs applicable for continuity check message (CCM):
  - End TLV
  - Sender ID TLV
  - Port Status TLV
  - Interface Status TLV
  - Organization-Specific TLV
- TLVs applicable for loopback message (LBM):
  - End TLV
  - Sender ID TLV
  - Data TLV
  - Organization-Specific TLV
- TLVs applicable for loopback reply (LBR):
  - End TLV
  - Sender ID TLV
  - Data TLV
  - Organization-Specific TLV
- TLVs applicable for linktrace message (LTM):
  - End TLV
  - LTM Egress Identifier TLV
  - Sender ID TLV
  - Organization-Specific TLV
- TLVs applicable for linktrace reply (LTR):
  - End TLV
  - LTR Egress Identifier TLV
  - Reply Ingress TLV
  - Reply Egress TLV
  - Sender ID TLV
  - Organization-Specific TLV

The following TLVs are currently supported in the applicable CFM PDUs:

- End TLV
- Reply Ingress TLV

- Reply Egress TLV
- LTR Egress Identifier TLV
- LTM Egress Identifier TLV
- Data TLV

Support for Additional Optional TLVs

The following additional optional TLVs are supported:

- Port Status TLV
- Interface Status TLV

MX Series routers support configuration of port status TLV and interface status TLV. Configuring the Port Status TLV allows the operator to control the transmission of the Port Status TLV in CFM PDUs.

NOTE: **Although Port Status TLV configuration statements are visible in the CLI on M120 and M320 routers, Port Status TLV cannot be configured on these systems. Port Status TLV can be enabled on a MEP interface only if it is a bridge logical interface, which is not possible on these systems.**

For configuration information, see the following sections:

- Port Status TLV on page 638
- Interface Status TLV on page 641

Port Status TLV

The Port Status TLV indicates the ability of the bridge port on which the transmitting MEP resides to pass ordinary data, regardless of the status of the MAC. The value of this TLV is driven by the MEP variable **enableRmepDefect**, as shown in Table 67 on page 639. The format of this TLV is shown in Table 66 on page 638.

Any change in the Port Status TLVs value triggers one extra transmission of that bridge ports MEP CCMs.

Table 66: Port Status TLV Format

Parameter	Octet (Sequence)
Type = 2	1
Length	2–3
Value (See Table 67 on page 639)	4

Table 67: Port Status TLV Values

Mnemonic	Ordinary Data Passing Freely Through the Port	Value
psBlocked	No: **enableRmepDefect** = false	1
psUp	Yes: **enableRmepDefect** = true	2

The MEP variable **enableRmepDefect** is a boolean variable indicating whether frames on the service instance monitored by the maintenance associations if this MEP are enabled to pass through this bridge port by the Spanning Tree Protocol and VLAN topology management. It is set to TRUE if:

- The bridge port is set in a state where the traffic can pass through it.
- The bridge port is running multiple instances of the spanning tree.
- The MEP interface is not associated with a bridging domain.

Configuring Port Status TLV

Junos OS provides configuration support for the Port Status TLV, allowing you to control the transmission of this TLV in CCM PDUs. The Junos OS provides this configuration at the continuity-check level. By default, the CCM does not include the Port Status TLV. To configure the Port Status TLV, use the **port-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check]** hierarchy level.

NOTE: Port Status TLV configuration is not mandated by IEEE 802.1ag. The Junos OS provides it in order to give more flexibility to the operator; however it receives and processes CCMs with a Port Status TLV, regardless of this configuration.

An example of the configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
        level number;
          maintenance-association identifier {
            continuity-check {
              interval number,
              loss-threshold number;
              hold-interval number;
              port-status-tlv; # Sets Port Status TLV
            }
          }
        }
      }
    }
```

```
  }
}
```

You cannot enable Port Status TLV transmission in the following two cases:

- If the MEP interface under the maintenance-association is not of type bridge.
- If the MEP is configured on a physical interface.

Displaying the Received Port Status TLV

The Junos OS saves the last received Port Status TLV from a remote MEP. If the received Port Status value does not correspond to one of the standard values listed in Table 67 on page 639, then the **show** command displays it as "unknown." You can display the last saved received Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***identifier*** **maintenance-association** ***identifier*** **local-mep** ***identifier*** **remote-mep** ***identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none # RX PORT STATUS
    Interface status TLV: none
```

Displaying the Transmitted Port Status TLV

The Junos OS saves the last transmitted Port Status TLV from a local MEP. If the transmission of the Port Status TLV has not been enabled, then the **show** command displays "none." You can display the last saved transmitted Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***identifier*** **maintenance-association** ***identifier*** **local-mep** ***identifier*** **remote-mep** ***identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up # TX PORT STATUS
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
```



```
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
```

Interface Status TLV

The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the IETF RFC 2863 IF-MIB. The format of this TLV is shown in Table 68 on page 641. The enumerated values are shown in Table 69 on page 641.

Table 68: Interface Status TLV Format

Parameter	Octet (Sequence)
Type = 4	1
Length	2–3
Value (See Table 69 on page 641)	4

Table 69: Interface Status TLV Values

Mnemonic	Interface Status	Value
isUp	up	1
isDown	down	2
isTesting	testing	3
isUnknown	unknown	4
isDormant	dormant	5
isNotPresent	notPresent	6
isLowerLayerDown	lowerLayerDown	7

NOTE: When the operational status of a logical interface changes from the down state (status value of 2) to the lower layer down state (status value of 7) and vice versa, the LinkDown SNMP trap is not generated. For example, if you configure an aggregated Ethernet interface bundle with a VLAN tag and add a physical interface that is in the operationally down state to the bundle, the operational status of the aggregated Ethernet logical interface bundle at that point is lower layer down (7). If you take the MIC associated with the interface offline, the LinkDown trap is not generated when the logical interface shifts from the lower layer down state to the down state.

Similarly, consider another sample scenario in which an physical interface is added to an aggregated Ethernet bundle that has VLAN tagging and the aggregated Ethernet logical interface is disabled. When the logical interface is disabled, the operational status of the logical interface changes to down. If you disable the physical interface that is part of the aggregated Ethernet bundle, the operational status of the aggregated Ethernet logical interface remains down. If you reenable the aggregated Ethernet logical interface, the operational status of it changes from down to lower layer down. The LinkDown SNMP trap is not generated at this point.

Configuring Interface Status TLV

The Junos OS provides configuration support for the Interface Status TLV, thereby allowing operators to control the transmission of this TLV in CCM PDUs through configuration at the continuity-check level.

NOTE: This configuration is not mandated by IEEE 802.1ag; rather it is provided to give more flexibility to the operator. The Junos OS receives and processes CCMs with the Interface Status TLV, regardless of this configuration.

The interface status TLV configuration is shown below:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
          level number;
          maintenance-association identifier {
            continuity-check {
              interval number;
              loss-threshold number;
              hold-interval number;
              interface-status-tlv; # Sets the interface status TLV
            }
          }
        }
      }
```

```
        }
      }
    }
```

NOTE: The Junos OS supports transmission of only three out of seven possible values for the Interface Status TLV. The supported values are 1, 2, and 7. However, the Junos OS is capable of receiving any value for the Interface Status TLV.

Displaying the Received Interface Status TLV

The Junos OS saves the last received Interface Status TLV from the remote MEP. If the received Interface Status value does not correspond to one of the standard values listed in Table 68 on page 641, then the **show** command displays "unknown."

You can display this last saved Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain *md5* maintenance-association *ma5* local-mep *2001* remote-mep *1001***

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none # displays the Interface Status TLV state
```

Displaying the Transmitted Interface Status TLV

The Junos OS saves the last transmitted Interface Status TLV from a local MEP. If the transmission of Interface Status TLV has not been enabled, then the **show** command displays "none."

You can display the last transmitted Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain *md5* maintenance-association *ma5* local-mep *2001* remote-mep *1001***

```
Maintenance domain name: md5, Format: string, Level: 5
```

```
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

MAC Status Defects

The Junos OS provides MAC status defect information, indicating that one or more of the remote MEPs is reporting a failure in its Port Status TLV or Interface Status TLV. It indicates “yes” if either some remote MEP is reporting that its interface is not isUp (for example, at least one remote MEPs interface is unavailable), or if all remote MEPs are reporting a Port Status TLV that contains some value other than psUp (for example, all remote MEPs Bridge Ports are not forwarding data). There are two **show** commands you can use to view the MAC Status Defects indication.

Use the **mep-database** command to display MAC status defects:

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain *md6* maintenance-association *ma6***

```
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM            : no
    Erroneous CCM received                  : no
    Cross-connect CCM received              : no
    RDI sent by some MEP                    : no
    Some remote MEP's MAC in error state    : yes # MAC Status Defects yes/no
  Statistics:
    CCMs sent                               : 1658
    CCMs received out of sequence           : 0
    LBMs sent                               : 0
    Valid in-order LBRs received            : 0
    Valid out-of-order LBRs received        : 0
    LBRs received with corrupted data       : 0
    LBRs sent                               : 0
    LTMs sent                               : 0
    LTMs received                           : 0
    LTRs sent                               : 0
    LTRs received                           : 0
    Sequence number of next LTM request     : 0
    1DMs sent                               : 0
    Valid 1DMs received                     : 0
```



```
    Invalid 1DMs received                  : 0
    DMMs sent                              : 0
    DMRs sent                              : 0
    Valid DMRs received                    : 0
    Invalid DMRs received                  : 0
  Remote MEP count: 1
    Identifier     MAC address        State    Interface
      200      00:05:85:73:39:4a       ok    xe-5/0/0.0
```

Use the **interfaces** command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: up, Port status TLV: up
    MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                : no
      Erroneous CCM received                      : no
      Cross-connect CCM received                  : no
      RDI sent by some MEP                        : no
      Some remote MEP's MAC in error state        : yes # MAC Status Defects
yes/no
    Statistics:
      CCMs sent                                   : 1328
      CCMs received out of sequence               : 0
      LBMs sent                                   : 0
      Valid in-order LBRs received                : 0
      Valid out-of-order LBRs received            : 0
      LBRs received with corrupted data           : 0
      LBRs sent                                   : 0
      LTMs sent                                   : 0
      LTMs received                               : 0
      LTRs sent                                   : 0
      LTRs received                               : 0
      Sequence number of next LTM request         : 0
      1DMs sent                                   : 0
      Valid 1DMs received                         : 0
      Invalid 1DMs received                       : 0
      DMMs sent                                   : 0
      DMRs sent                                   : 0
      Valid DMRs received                         : 0
      Invalid DMRs received                       : 0
    Remote MEP count: 1
      Identifier     MAC address        State    Interface
        200      00:05:85:73:39:4a       ok    xe-5/0/0.0
```

Configuring Remote MEP Action Profile Support

Based on values of **interface-status-tlv** and **port-status-tlv** in the received CCM packets, a specific action, such as **interface-down**, can be taken using the **action-profile** options. Multiple action profiles can be configured on the router, but only one action profile can be assigned to a remote MEP.

The action profile can be configured with at least one event to trigger the action; but the action will be triggered if any one of these events occurs. It is not necessary for all of the configured events to occur to trigger **action**.

An action-profile can be applied only at the remote MEP level.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile tlv-action {
  event {
    # If interface status tlv with value specified in the config is received
    interface-status-tlv down|lower-layer-down;
    # If port status tlv with value specified in the config is received
    port-status-tlv blocked;
    # If connectivity is lost to the peer */
    adjacency-loss;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
  default-actions interface-down;
  }
# domains
  maintenance-domain identifier {
    # maintenance domain level (0-7)
    level number;
    # association
    maintenance-association identifier {
      mep identifier {
        interface ge-x/y/z.w;
        remote-mep identifier {
          # Apply the action-profile for the remote MEP
          action-profile tlv-action;
        }
      }
    }
  }
```

Monitoring a Remote MEP Action Profile

You can use the **show oam ethernet connectivity-fault-management mep-database** command to view the action profile status of a remote MEP, as in the following example:

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
  Auto-discovery: enabled, Priority: 0
```

```
 Interface status TLV: none, Port status TLV: none  # last status TLVs transmitted
by the router
 Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

 Remote MEP identifier: 200, State: ok # displays the remote MEP name and state

   MAC address: 00:05:85:73:96:1f, Type: Configured
   Interface: ge-1/0/8.0
   Last flapped: Never
   Remote defect indication: false
   Port status TLV: none
   Interface status TLV: lower-layer-down
   Action profile: juniper  # displays remote MEP’s action profile identifier
     Last event: Interface-status-tlv lower-layer-down  # last remote MEP event

                                                         # to trigger action
     Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
                                 # action occurrence time
```

Related Documentation

- **connectivity-fault-management on page 1113**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655

Configuring MAC Flush Message Processing in CET Mode

In carrier Ethernet transport (CET) mode, MX Series routers are used as provider edge (PE) routers, and Nokia Siemens Networks A2200 Carrier Ethernet Switches (referred to as E-domain devices) that run standard-based protocols are used in the access side. On the MX Series routers, VPLS pseudowires are configured dynamically through label distribution protocol (LDP). On the E-domain devices, topology changes are detected through connectivity fault management (CFM) sessions running between the E-domain devices and the MX Series PE routers. The MX Series PE routers can bring the carrier Ethernet interface down if there is CFM connectivity loss. This triggers a local MAC flush as well as a targeted label distribution protocol (T-LDP) MAC flush notification that gets sent towards the remote MX Series PEs to trigger MAC flush on them.

In CET inter-op mode, MX Series routers need to interoperate with the Nokia Siemens Networks Ax100 Carrier Ethernet access devices (referred to as A-domain devices) that run legacy protocols. Nokia Siemens Networks A4100 and A8100 devices act as an intermediate between the MX Series PE routers and A-domain devices. These intermediate devices perform interworking function (IWF) procedures so that operations administration management (OAM) sessions can be run between MX Series routers and A-domain devices. There are no VPLS pseudowires between the MX Series PE routers and the Nokia Siemens Networks A4100 and A8100 intermediate devices, so there is no LDP protocol running between the PE routers to send topology change notifications. In order to communicate topology changes, MX Series routers can trigger a MAC flush and propagate it in the core. MX Series routers can use action profiles based upon the connection protection type length value (TLV) event. The action profile brings down the carrier edge logical interface in MX Series PE routers, which will trigger a local MAC flush and also propagate the topology change to the core using LDP notification.

For VPLS there is no end-to-end connectivity monitored. The access rings are independently monitored by running CFM down multiple end points (MEPs) on the working and protection paths for each of the services between the E-domain devices and the MX Series PE routers, and between the A-domain devices and the MX Series PE routers the IWF hosted by the Nokia Siemens Networks A-4100 devices. When there is a connectivity failure on the working path, the Nokia Siemens Networks Ax200 devices perform a switchover to the protection path, triggering a topology change notification (in the form of TLVs carried in CCM) to be sent on the active path.

Figure 41: CET inter-op Dual Homed Topology

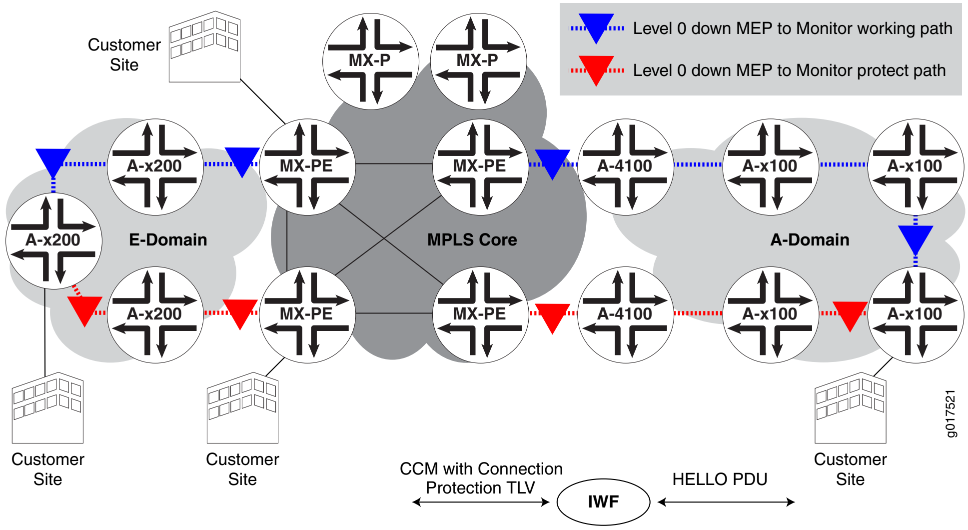

Figure 41 on page 649 describes the dual homed topology on MX Series PE routers connected to the A-domain. When an A-domain device triggers a switchover, it starts switching the service traffic to the new active path. This change is communicated in the HELLO protocol data units (PDUs) sent by that A-domain device on the working and protection paths. When the IWF in A4100 recieves these HELLO PDUs, it converts them to standard CCM messages and also inserts a connection protection TLV. The “Protection-in-use” field of the connection protection TLV is encoded with the currently active path, and is included in the CCM message. CCM messages are received by the MX Series PE routers through the VLAN spoke in A4100. In the above dual homed scenario, one MX Series PE router monitors the working path, and the other MX Series PE router monitors the protection path.

A MAC flush occurs when the CFM session that is monitoring the working path detects that the service traffic has moved to the protection path or when the CFM session that is monitoring the protection path detects that the service traffic has moved to the working path.

Figure 42: CET inter-op Dual Attached Topology

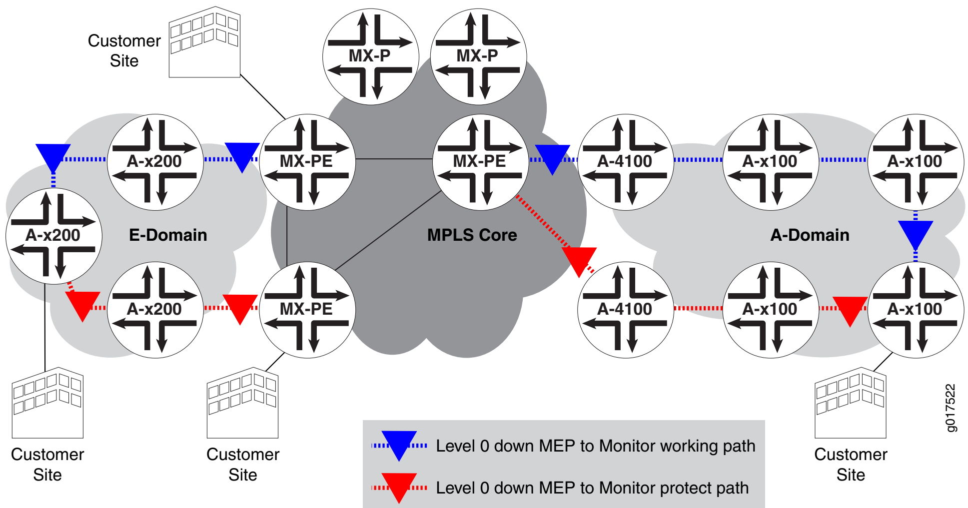

Figure 42 on page 650 describes the dual attached topology on MX Series PE routers connected to the A-domain. The MAC flush mechanism used in this case is also the same as the one used for the A-domain in the dual homed scenario (Figure 1). However in this case both the CFM sessions are hosted by only one MX Series PE router. When Ax100 in the A-domain detects topology changes, the MX Series PE router receives the connection protection TLV in the CCM message for the working and protection paths with the value of “Protection-in-use” indicating which path is the active one. Based upon the event that is generated for the CFM session, the MX Series PE router will bring down the appropriate interface which will trigger a local MAC flush.

Configuring a Connection Protection TLV Action Profile

An action profile can be configured to perform the **interface-down** action based on the values of **connection-protection-tlv** in the received CCM packets.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
    # If a connection protection TLV with a “Protection-in-use” value of SET is received */
    connection-protection-tlv <using-protection-path>;
    # If a connection protection TLV with a “Protection-in-use” value of RESET is received
      */
    connection-protection-tlv <using-working-path>;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
}
```


Related Documentation

- *connection-protection-tlv*

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Creating a Maintenance Association on page 610
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- Configuring Rate Limiting of Ethernet OAM Messages on page 655

Example: Configuring an Action Profile Based on Connection Protection TLVs

This example shows how to configure an action profile based on the connection protection TLV for the purposes of triggering MAC flushes based on topology changes in a CET network.

- Requirements on page 651
- Overview and Topology on page 651
- Configuration on page 652

Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.2 or later
- A MX series PE router

Overview and Topology

The physical topology of a CET network using MX series PE routers is shown in Figure 43 on page 652

Figure 43: Topology of CET network

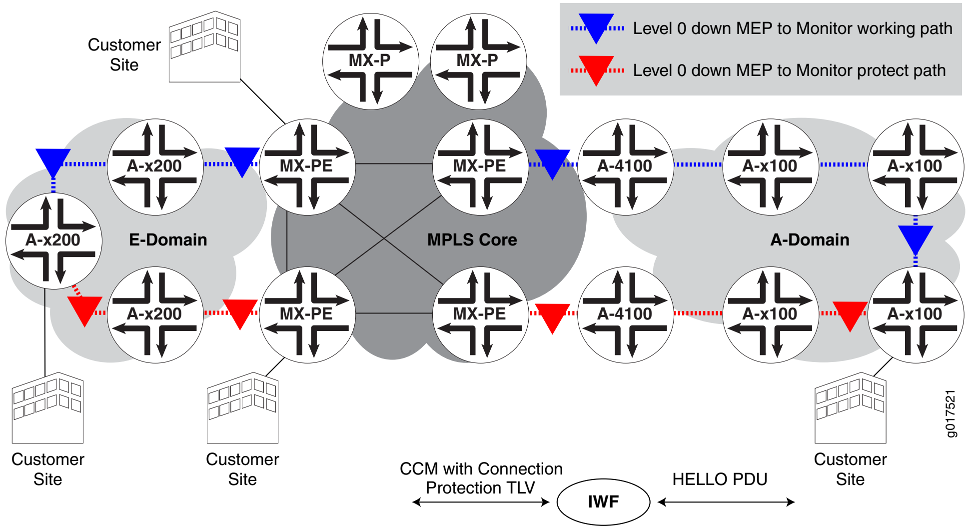

The following definitions describe the meaning of the device abbreviation and terms used in Figure 43 on page 652.

- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.
- E-domain—Nokia Siemens Networks Carrier Ethernet Switches that run standard based protocols and are used in the access side.
- A-domain—Nokia Siemens Networks Carrier Ethernet Switches that run legacy protocols.

Configuration

Step-by-Step Procedure

To configure an action profile based on the connection protection TLV, preform these tasks:

1. Configure an action profile

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
```

2. If the connection protection TLV is received with a “Protection-in-use” value of SET, then the connection protection TLV should use the protection path

```
connection-protection-tlv <using-protection-path>;
```

3. If the connection protection TLV is received with a “Protection-in-use” value of RESET, then the connection protection TLV should use the working path

```
connection-protection-tlv <using-working-path>;
}
```

4. Configure the action profile to bring the interface down

```
action {
  /* Bring the interface down */
  interface-down;
}
}
```

Results Check the results of the configuration

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
    connection-protection-tlv <using-protection-path>;
    connection-protection-tlv <using-working-path>;
  }
  action {
    interface-down;
  }
}
```

Related Documentation

- *connection-protection-tlv*
- Configuring MAC Flush Message Processing in CET Mode on page 648

Configuring M120 and MX Series Routers for CCC Encapsulated Packets

IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets Overview

Layer 2 virtual private network (L2VPN) is a type of virtual private network service used to transport customer's private Layer 2 traffic (for example, Ethernet, ATM or Frame Relay) over the service provider's shared IP/MPLS infrastructure. The service provider edge (PE) router must have an interface with circuit cross-connect (CCC) encapsulation to switch the customer edge (CE) traffic to the public network.

The IEEE 802.1ag Ethernet Connectivity Fault Management (CFM) is an OAM standard used to perform fault detection, isolation, and verification on virtual bridge LANs. M120 and MX Series routers provide CFM support for bridge/VPLS/routed interfaces and support 802.1ag Ethernet OAM for CCC encapsulated packets.

CFM Features Supported on Layer 2 VPN Circuits

CFM features supported on L2VPN circuits are as follows:

- Creation of up/down MEPs at any level on the CE-facing logical interfaces.
- Creation of MIPs at any level on the CE-facing logical interfaces.

- Support for continuity check, loopback, and linkrace protocol.
- Support for the Y1731 Ethernet Delay measurement protocol.
- Support for action profiles to bring the CE-facing logical interfaces down when loss of connectivity is detected.

Figure 44: Layer 2 VPN Topology

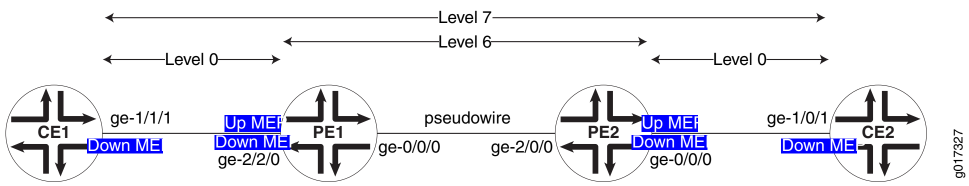

To monitor the L2VPN circuit, a CFM up MEP (Level 6 in Figure 44 on page 654) can be configured on the CE-facing logical interfaces of provider edge routers PE1 and PE2. To monitor the CE-PE attachment circuit, a CFM down MEP can be configured on the customer logical interfaces of CE1-PE1 and CE2-PE2 (Level 0 in Figure 44 on page 654).

Configuring CFM for CCC Encapsulated Packets

The only change from the existing CLI configuration is the introduction of a new command to create a MIP on the CE-facing interface of the PE router.

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        # Define a maintenance domains for each default level.
        #; These names are specified as DEFAULT_level_number
        maintenance-domain DEFAULT_x {
          # L2VPN CE interface
          interface (ge | xe)-fpc/pic/port.domain;
        }
        {
        level number;
        maintenance-association identifier {
          mep mep-id {
            direction (up | down);
            # L2 VPN CE interface on which encapsulation family CCC is configured.
            interface (ge | xe)-fpc/pic/port.domain;
            auto-discovery;
            priority number;
          }
        }
      }
    }
  }
  }
}
```

See Also

- connectivity-fault-management on page 1113
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

Configuring Rate Limiting of Ethernet OAM Messages

The M320 with Enhanced III FPC, M120, M7i, M10 with CFEB, and MX Series routers support rate limiting of Ethernet OAM messages. Depending on the connectivity fault management (CFM) configuration, CFM packets are discarded, sent to the CPU for processing, or flooded to other bridge interfaces. This feature allows the router to intercept incoming CFM packets for prevention of DoS attacks.

You can apply rate limiting of Ethernet OAM messages at either of two CFM policing levels, as follows:

- Global-level CFM policing—uses a policer at the global level to police the CFM traffic belonging to all the sessions.
- Session-level CFM policing—uses a policer created to police the CFM traffic belonging to one session.

To configure global-level CFM policing, include the **policer** statement and its options at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

To configure session-level CFM policing, include the **policer** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* level *number* maintenance-association *ma-name*]** hierarchy level.

The following example shows a CFM policer used for rate-limiting CFM:

```
[edit]
firewall {
  policer cfm-policer {
    if-exceeding {
      bandwidth-limit 8k;
      burst-size-limit 2k;
    }
```

```
      then discard;
    }
  }
```

Case 1: Global-Level CFM Policing

This example shows a global level policer, at the CFM level, for rate-limiting CFM. The **continuity-check** ***cfm-policer*** statement at the global **[edit protocols oam ethernet connectivity-fault-management policer]** hierarchy level specifies the policer to use for policing all continuity check packets of the CFM traffic belonging to all sessions. The **other** ***cfm-policer1*** statement at the **[edit protocols oam ethernet connectivity-fault-management policer]** hierarchy level specifies the policer to use for policing all non-continuity check packets of the CFM traffic belonging to all sessions. The **all** ***cfm-policer2*** statement specifies to police all CFM packets with the specified policer *cfm-policer2*. If the **all** ***policer-name*** option is used, then the user cannot specify the previous **continuity-check** and **other** options.

```
[edit protocols oam ethernet]
connectivity-fault-management {
  policer {
    continuity-check cfm-policer;
    other cfm-policer1 ;
    all cfm-policer2;
    }
  }
```

Case 2: Session-Level CFM Policing

This example shows a session-level CFM policer used for rate-limiting CFM. The **policer** statement at the session **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*****]** hierarchy level specifies the policer to use for policing only continuity check packets of the CFM traffic belonging to the specified session. The **other** ***cfm-policer1*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*****]** hierarchy level specifies the policer to use for policing all non-continuity check packets of the CFM traffic belonging to this session only. The **all** ***cfm-policer2*** statement specifies to police all CFM packets with the specified policer *cfm-policer2*. If the **all** ***policer-name*** option is used, then the user cannot specify the previous **continuity-check** and **other** options.

```
[edit protocols oam ethernet]
connectivity-fault-management {
  maintenance-domain md {
    level number;
    maintenance-association ma {
      continuity-check {
        interval 1s;
      }
      policer {
        continuity-check cfm-policer;
        other cfm-policer1;
        all cfm-policer2;
        }
      }
      mep 1 {
        interface ge-3/3/0.0;
        direction up;
        auto-discovery;
      }
```



```
    }
  }
```

In the case of global CFM policing, the same policer is shared across multiple CFM sessions. In per-session CFM policing, a separate policer must be created to rate-limit packets specific to that session.

NOTE:
Service-level policer configuration for any two CFM sessions on the same interface at different levels must satisfy the following constraints if the direction of the sessions is the same:

- If one session is configured with **policer all**, then the other session cannot have a **policer all** or **policer other** configuration.
- If one session is configured with **policer other**, then the other session cannot have a **policer all** or **policer other** configuration.

A commit error will occur if such a configuration is committed.

NOTE: Policers with PBB and MIPs are not supported.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Domain on page 603
- Configuring Maintenance Intermediate Points (MIPs) on page 604
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Linktrace Protocol in CFM on page 626
- Configuring Ethernet Local Management Interface on page 627
- Configuring Port Status TLV and Interface Status TLV on page 635
- Configuring MAC Flush Message Processing in CET Mode on page 648
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 653
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Unified ISSU for 802.1ag CFM

A unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is automatically enabled for the Connectivity Fault

Management (CFM) protocols and interoperates between local and remote maintenance endpoints (MEPs).

The Junos OS provides support for unified ISSU using the loss threshold type length value (TLV), which is automatically enabled for CFM. TLVs are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and optional information in a protocol data unit (PDU). The loss threshold TLV indicates the loss threshold value of a remote MEP. The loss threshold TLV is transmitted as part of the CFM continuity check messages.

NOTE: **Starting in Junos OS Release 15.1, configuring ISSU with CFM (802.1ag) is supported only on MX and PTX routers that support TLV. Interoperation with other vendors is not supported.**

During a unified ISSU, the control plane may go down for several seconds and cause CFM continuity check packets to get dropped. This may cause the remote MEP to detect a connectivity loss and mark the MEP as down. To keep the MEP active during a unified ISSU, the loss threshold TLV communicates the minimum threshold value the receiving MEP requires to keep the MEP active. The receiving MEP parses the TLV and updates the loss threshold value, but only if the new threshold value is greater than the locally configured threshold value.

An overview of CFM is described starting in "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 596, and you should further observe the additional requirements described in this topic.

Table 70 on page 658 shows the Loss Threshold TLV format.

Table 70: Loss Threshold TLV Format

Parameter	Octet (sequence)	Description
Type=31	1	Required. Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.
Length=12	2	Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field.
OUI	3	Optional. Organization unique identifier (OUI), which is controlled by the IEEE and is typically the first three bytes of a MAC address (Juniper OUI 0x009069).
Subtype	1	Optional. Organizationally defined subtype.
Value	4	Optional. Loss threshold value.
Flag	4	Optional. Bit0 (identifies an ISSU is in progress) Bit1-31 (reserved)

Junos OS provides configuration support for the **convey-loss-threshold** statement, allowing you to control the transmission of the loss threshold TLV in continuity check messages PDUs. The **convey-loss-threshold** statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the **convey-loss-threshold** statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress. The Junos OS provides this configuration at the continuity-check level. By default, continuity check messages do not include the loss threshold TLV.

To configure the convey loss threshold, use the **convey-loss-threshold** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check]** hierarchy level.

For the remote MEP, the loss threshold TLV is transmitted only during the unified ISSU if the **convey-loss-threshold** statement is not configured. The remote MEP switches back to the default loss threshold if no loss threshold TLV is received or the TLV has a default threshold value of 3.

An example of the ISSU configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
        level number;
          maintenance-association identifier {
            continuity-check {
              convey-loss-threshold;
              interval number;
              loss-threshold number;
              hold-interval number;
            }
          }
        }
      }
    }
  }
}
```

The Junos OS saves the last received loss threshold TLV from the remote MEP. You can display the last saved loss threshold TLV that is received by the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md3 maintenance-association ma5 local-mep 2 remote-mep 1
Maintenance domain name: md3, Format: string, Level: 3
  Maintenance association name: ma3, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
```

```
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
    Loss Threshold TLV:
      Loss Threshold: 3  , Flag:    0x0

 Remote MEP identifier: 1, State: ok
    MAC address: 00:1f:12:b7:ce:79, Type: Learned
    Interface: xe-4/1/1.0
    Last flapped: Never
    Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
    Effective loss threshold: 3 frames
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
    Connection Protection TLV:
      Prefer me: no, Protection in use: no, FRR Flag: no
    Loss Threshold TLV:    #Displays last received value
      Loss Threshold: 3  , Flag:    0x0
```

The Junos OS saves the last transmitted loss threshold TLV from a local MEP. You can display the last transmitted loss threshold TLV and the effective loss (operational) threshold for the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***identifier*** **maintenance-association** ***identifier*** **local-mep** ***identifier*** **remote-mep** ***identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md3 maintenance-association ma5 local-mep 2 remote-mep 1
Maintenance domain name: md3, Format: string, Level: 3
  Maintenance association name: ma3, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
    Loss Threshold TLV:    #Displays last transmitted value
      Loss Threshold: 3  , Flag:    0x0

   Remote MEP identifier: 1, State: ok
    MAC address: 00:1f:12:b7:ce:79, Type: Learned
    Interface: xe-4/1/1.0
    Last flapped: Never
    Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
    Effective loss threshold: 3 frames    #Displays operational threshold
    Remote defect indication: falsePort status TLV: none
    Interface status TLV: none
    Connection Protection TLV:
      Prefer me: no, Protection in use: no, FRR Flag: no
    Loss Threshold TLV:
      Loss Threshold: 3  , Flag:    0x0
```


Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, configuring ISSU with CFM (802.1ag) is supported only on MX and PTX routers that support TLV.

Related Documentation

- Example: Configuring Ethernet CFM over VPLS on page 675
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

Configuring Continuity Check Messages for Better Scalability

This topic describes how to configure CCM for better scalability. Junos OS provides enhancements to trigger faster protection-switching and convergence in the event of failures in Ethernet domains for Carrier Ethernet services. These enhancements can be used when CE devices in the Ethernet domain detect faster service failures and propagates the information in the interface-status TLV of the continuity-check messages (CCMs). When CCMs are received, PE devices can perform certain actions which facilitates faster protection-switching and convergence.

To configure CCM for better scalability:

- You can apply an action profile to provide faster protection switching for point-to-point network topologies with local switching configured. See “Configuring Faster Protection Switching for Point-to-Point Network Topologies” on page 662.
- You can apply an action profile to provide faster convergence for dual-homed multipoint-to-multipoint network topologies. See “Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies” on page 663.
- You can assign a primary virtual LAN (VLAN) ID in the maintenance association for increased flexibility in the number of tags. See “Configuring a Primary VLAN ID for Increased Flexibility” on page 664.
- You can configure a maintenance association to accept a different maintenance association identifier (ID) from a neighbor by including a **remote-maintenance-association** statement. See “Configuring a Remote Maintenance Association to Accept a Different ID” on page 665.

Related Documentation

- Configuring Faster Protection Switching for Point-to-Point Network Topologies on page 662
- Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies on page 663
- Configuring a Primary VLAN ID for Increased Flexibility on page 664
- Configuring a Remote Maintenance Association to Accept a Different ID on page 665

Configuring Faster Protection Switching for Point-to-Point Network Topologies

You can apply an action profile to provide faster protection switching for point-to-point network topologies with local switching configured. In a normal state, CCM sessions are configured on the working and protect interfaces. The CCM packets transmitted contain an interface-status TLV with the value up on the working interface and value down on the protect interface. When a link fails on the working interface, the protect interface starts receiving the interface-status TLV as up. With the profile configuration, if the interface-status TLV received on the protect interface is up, the working interface is automatically marked as **interface-down**.

To configure the **interface-status-tlv** down event, include the **interface-status-tlv down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure **interface-down** as the action profile's action, include the **interface-down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* action]** hierarchy level.

To configure **interface-down *peer-interface*** as the clear-action, include **interface-down *peer-interface*** at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* clear-action]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile p1 {
      event {
        interface-status-tlv down;
      }
      action {
        interface-down;
      }
      clear-action {
        interface-down peer-interface;
      }
    }
  }
}
```

In this action profile configuration, when the interface-status TLV is received as up, the *peer-interface* is marked as down.

The *peer-interface* is configured in the **protect-maintenance-association** statement. Consider the following example using the **protect-maintenance-association** statement in the configuration:

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile p1 {
      event {
        adjacency-loss;
```



```
          }
          action {
            interface-down;
          }
          clear-action {
            interface-down peer-interface;
          }
        }
        maintenance-domain nsn {
          level 5;
          maintenance-association ma1 {
            protect-maintenance-association ma2;
            continuity-check {
              interval 100ms;
              connection-protection-tlv;
            }
            mep 100 {
              interface ge-1/1/0.0;
              direction down;
              auto-discovery;
            }
          }
          maintenance-association ma2 {
            continuity-check {
              interval 100ms;
              connection-protection-tlv;
          }
          mep 101 {
              interface ge-1/2/0.0;
              direction down;
              auto-discovery;
          }
          remote-mep 100
            action-profile p1;
          }
        }
      }
    }
```

Related Documentation

- connectivity-fault-management on page 1113
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625

Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies

You can apply an action profile to provide faster convergence for dual-homed multipoint-to-multipoint network topologies. If a multipoint-to-multipoint Ethernet service uses MAC-based forwarding and stale MAC addresses exist in the learning tables, this can result in traffic black holes in the network where incoming traffic is silently discarded, without informing the source that the data did not reach its intended recipient. With the profile configuration, if the interface-status TLV received on the protect interface is up, then the interface-status TLV on the working interface is marked as down and the PE device for the protect interface propagates a remote MAC-flush message to the PE devices in the virtual private LAN service (VPLS) by using TLDP-MAC-FLUSH. The MAC flush avoids traffic blackholing due to stale mac-db entries.

To configure the **interface-status-tlv** down event, include the **interface-status-tlv down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure **propagate-remote-flush** as the action profile's action, include the **propagate-remote-flush** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* action]** hierarchy level.

To configure **propagate-remote-flush** as the clear-action, include the **propagate-remote-flush** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* clear-action]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile test {
      event {
        interface-status-tlv down;
      }
      action {
        propagate-remote-mac-flush;
      }
      clear-action {
        propagate-remote-mac-flush;
      }
    }
  }
}
```

In this action profile configuration, when the incoming CCM packet contains the interface-status TLV with value down, the **propagate-remote-mac-flush** action is triggered for the action-profile.

Related Documentation

- Configuring MAC Flush Message Processing in CET Mode on page 648
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- **connectivity-fault-management on page 1113**
- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625

Configuring a Primary VLAN ID for Increased Flexibility

You can assign a primary virtual LAN (VLAN) ID in the maintenance association for increased flexibility in the number of tags. When a **vlan-range** or **vlan-id-list** is configured on an interface, the service OAM must run on one of the VLANs. The VLAN assigned for service monitoring is considered the primary VLAN. If a **primary-vid** is not configured, Junos OS assigns the first VLAN from the **vlan-range** or **vlan-id-list**. In earlier releases, Junos OS assigned VLAN 4095.

To configure a primary VLAN ID, you can specify the **primary-vid** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
maintenance domain md3 {
  level 3;
  maintenance-association ma3 {
    primary-vid 2000;
    continuity-check {
      interval 10ms;
      connection-protection-tlv;
    }
    mep 2 {
      interface ge-2/2/0.0;
      direction up;
      auto-discovery;
    }
  }
}
```

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- *connection-protection-tlv*
- Creating a Maintenance Association on page 610
- **connectivity-fault-management on page 1113**

Configuring a Remote Maintenance Association to Accept a Different ID

You can configure a maintenance association to accept a different maintenance association identifier (ID) from a neighbor by including a **remote-maintenance-association** statement. The 802.1ag CCM sessions expect the same maintenance association identifier from its neighbors. If there is a maintenance association identifier mismatch, the PDUs are marked as error PDUs. If a **remote-maintenance-association** statement is configured, a different maintenance association identifier is accepted and the 802.1ag CCM sessions do not mark the CCM PDUs as error PDUs when the maintenance-association name is the same as the name specified in the **remote-maintenance-association** statement.

To configure a remote maintenance association, include the **remote-maintenance-association** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
maintenance domain md3 {
  level 1;
  maintenance-association ma3 {
    remote-maintenance-association fix-ma;
```

```
            continuity-check {
                interval 10ms;
                connection-protection-tlv;
            }
            mep 2 {
                interval ge-2/2/0.0;
                direction up;
                auto-discovery;
            }
        }
    }
```

Using this configuration, interoperability is improved for CCMs with low-end CE devices supporting fixed maintenance association identifier configurations.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Creating a Maintenance Association on page 610
- **connectivity-fault-management on page 1113**
- *connection-protection-tlv*

Enabling Enhanced Connectivity Fault Management Mode

You can enable enhanced connectivity fault management (CFM) mode to enable effective Ethernet OAM deployment in scaling networks. On enabling enhanced CFM mode, Junos OS supports 32, 000 maintenance association end points (MEPs) and maintenance intermediate points (MIPs) each per chassis for bridge, VPLS, L2VPN, and CCC domains. In previous releases, Junos OS supports 8, 000 MEPs and 8000 MIPS per chassis. If you do not enable enhanced CFM, Junos OS continues to support existing number of MIPs and MEPs per chassis.

NOTE: To support enhanced CFM mode, configure the network services mode on the router as **enhanced-ip**. If the network services mode is not **enhanced-ip**, and you have enabled enhanced CFM, the following warning message is displayed:
[edit protocols oam ethernet]
'connectivity-fault-management'
enhanced ip is not effective please configure enhanced ip and give router reboot

To enable enhanced CFM mode, perform the following steps:

1. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

   ```
   [edit]
   user@host# edit protocols oam ethernet connectivity-fault-management
   ```

2. Enable effective Ethernet OAM deployment by enabling enhanced CFM mode.


```
[edit protocols oam ethernet connectivity-fault-management ]
user@host# set enhanced-cfm-mode
```

3. Commit the mode change. A warning message is displayed asking you to restart CFM. If you do not restart CFM, CFM is automatically restarted by Junos OS.

```
[edit protocols oam ethernet connectivity-fault-management ]
user@host # commit
[edit protocols oam ethernet]
'connectivity fault management'
CFM mode change is catastrophic. cfmd will be restarted
commit complete
```

4. To verify if the enhanced CFM mode has been configured, use the **show oam ethernet connectivity-fault-management state** command.

```
[edit protocols oam ethernet connectivity-fault-management ]
user@host# run show oam ethernet connectivity-fault-management
 enhanced-cfm-mode;
 traceoptions {
     file cfmd.log size 1g;
 }
 maintenance-domain md6 {
     level 6;
     maintenance-association ma6 {
         continuity-check {
             interval 1s;
         }
         mep 102 {
             interface ge-0/0/0.0;
             direction up;
         }
     }
 }
```

Related Documentation

- enhanced-cfm-mode on page 1011

Understanding CFM Monitoring between CE and PE Devices

You can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device. When the interface is down, CFM propagates the status of the interface in the CC messages. The CC message informs the customer edge device that the provider edge device is down.

You can configure CFM monitoring using either of the following two options:

- Interface Status TLV (Type, Length, and Value)—You can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device by using Interface Status TLV. When the interface is down, CFM propagates the status of the interface using interface status TLV. The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the

IETF RFC 2863 IF-MIB. Thus, the customer edge device is aware that the provider edge device is down. To configure CFM monitoring using Interface Status TLV, use the **interface-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain* maintenance-association *maintenance-association* continuity-check** hierarchy level. This is the standard option.

- RDI (Remote Defect Indication)—Starting in Junos OS Release 17.3R1, you can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device by using the remote defect indication (RDI) bit. When you enable CFM monitoring, CFM propagates the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages. Thus, the customer edge device is aware that the provider edge device is down. The RDI bit is cleared when the service is back up. To configure CFM monitoring using the RDI bit, use the **interface-status-send-rdi** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain* maintenance-association *maintenance-association* continuity-check** hierarchy level. This option is required if the customer edge device does not support Interface Status TLV.

NOTE: **When the interface is set to CCC down and you have configured RDI, then RDI bit is sent. CFM does not monitor the status of the interface. If CCC down is set when the interface is not standby, RDI bit is sent with the CC messages if you have configured RDI.**

Single Active Multi-homing Use Case using RDI bit

Consider the following topology where there are two provider edge devices (PE1 and PE2) as well as two customer edge devices (CE1 and CE2). PE1 is in active state while PE2 is in standby state. CFM down MEP is configured between the PE and CE. CFM detects that the CCC down and because CFM down MEP is configured, the CC messages generated have the RDI bit. The CC messages from PE2 to CE2 have the RDI bit set to indicate the blocked state. When PE2 becomes active, CCM down is cleared and the RDI bit is cleared from the subsequent CC messages.

Active/Active Multihoming Use case using RDI bit

Consider the topology where there are two provider edge devices (PE1 and PE2) and two customer edge devices (CE1 and CE2). PE1 is in active state while PE2 is in standby state. If CFM down MEP is not configured between the PE and CE to monitor the link connectivity, the CC messages generated do not have the RDI bit. CFM down MEP is configured between the PE and CE. CFM detects that the CCC down and because CFM down MEP is configured, the CC messages generated have the RDI bit. The CC messages from PE2 to CE2 have the RDI bit set to indicate the blocked state. When PE2 becomes active, CCM down is cleared and the RDI bit is cleared from the subsequent CC messages.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, you can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device by using the remote defect indication (RDI) bit.

Related Documentation

- interface-status-tlv on page 1018
- Configuring Port Status TLV and Interface Status TLV on page 635
- interface-status-send-rdi on page 1019

Example: Configuring Ethernet CFM on Physical Interfaces

This example shows the configuration of Ethernet connectivity fault management (CFM) on physical interfaces.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.3 or later.

Overview

CFM can be used to monitor the physical link between two routers. This functionality is similar to that supported by the IEEE 802.3ah LFM protocol.

In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces. On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces.

NOTE: **The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.**

Configuration

In the following example, two routers (Router 1 and Router 2) are connected by a point-to-point Gigabit Ethernet link. The link between these two routers is monitored using CFM. This is shown in Figure 45 on page 670. The single boundary is a “down mep” in CFM terminology.

Figure 45: Ethernet CFM on Physical Interfaces

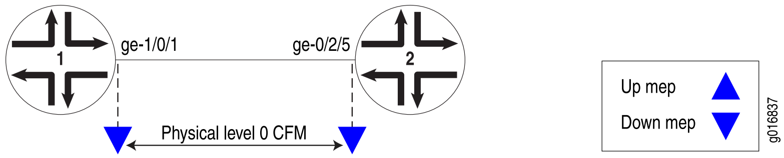

To configure Ethernet CFM on physical interfaces, perform these tasks:

CLI Quick Configuration Router 1

Configure the interface and CFM:

```
[edit]
interfaces ge-1/0/1 {
  unit 0 {
    family inet;
  }
}

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

The configuration on Router 2 mirrors that on Router 1, with the exception of the *mep-id*.

Router 2

Configure the interface and CFM:

```
[edit]
interfaces ge-0/2/5 {
  unit 0 {
    family inet;
  }
}

protocols {
```

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

To verify that the physical interface is configured correctly for CFM, use the **show interface** command. To verify the CFM configuration, use one or more of the **show oam ethernet connectivity-fault-management** commands listed in the CLI Explorer.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Operations, Administration, and Maintenance on page 594
- Ethernet OAM Connectivity Fault Management on page 595
- Example: Configuring Ethernet CFM over VPLS on page 675
- Example: Configuring Ethernet CFM on Bridge Connections on page 671

Example: Configuring Ethernet CFM on Bridge Connections

In this example, both the customer and service provider are running Ethernet CFM over a simple bridge network. The network is shown in Figure 46 on page 672. The customer has configured Ethernet CFM on MX Series routers L2-CE1 and L2-CE2. The service provider has configured Ethernet CFM on MX Series routers PE1 and PE2.

NOTE: **The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.**

The service provider is using CFM level 3 for the link between PE1 and PE2 and level 5 from one CE facing port to the other. The customer is using CFM level 7. The boundaries are marked with “up mep” and “down mep” CFM terminology in the figure.

Figure 46: Ethernet CFM over a Bridge Network

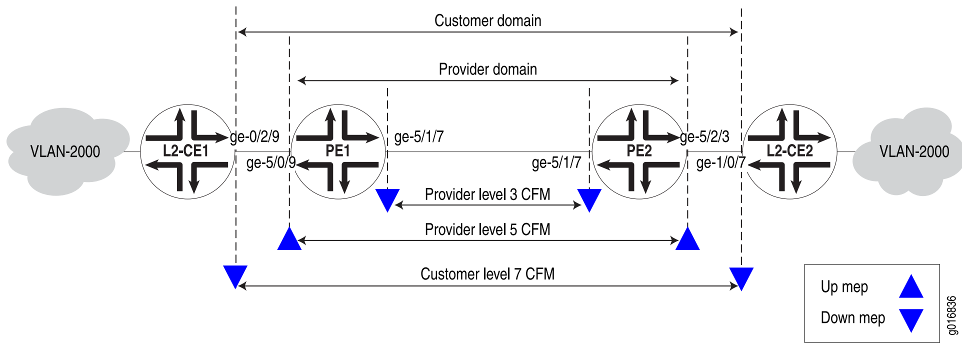

Here are the configurations of CFM on the customer routers.

CFM on L2-CE1

```
[edit interfaces]
ge-0/2/9 {
   vlan-tagging;
   unit 0 {
      vlan-id 2000;
   }
}

[edit protocols oam ethernet]
connectivity-fault-management {
   maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
         continuity-check {
            interval 1s;
         }
         mep 700 {
            interface ge-0/2/9.0;
            direction down;
            auto-discovery;
         }
      }
   }
}
```

CFM on L2-CE2

```
[edit interfaces]
ge-1/0/7 {
   vlan-tagging;
   unit 0 {
      vlan-id 2000;
   }
}

[edit protocols oam ethernet]
connectivity-fault-management {
   maintenance-domain customer {
      level 7;
      maintenance-association customer-site2 {
```



```
            continuity-check {
                interval 1s;
            }
            mep 800 {
                interval ge-1/0/7.0;
                direction down;
                auto-discovery;
            }
        }
    }
}
```

Here are the configurations of CFM on the provider routers.

CFM on PE1

```
[edit interfaces]
ge-5/0/9 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}
ge-5/1/7 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}

[edit bridge-domains]
bridge-vlan2000 {
    domain-type bridge;
    vlan-id 2000;
    interface ge-5/0/9.0;
    interface ge-5/1/7.0;
}

[edit protocols oam ethernet connectivity-fault-management]
maintenance-domain provider-outer {
    level 5;
    maintenance-association provider-outer-site1 {
        continuity-check {
            interval 1s;
        }
        mep 200 {
            interface ge-5/0/9.0;
            direction up;
            auto-discovery;
        }
    }
}
maintenance-domain provider-inner {
```

```
    level 3;
    maintenance-association provider-inner-site1 {
        continuity-check {
            interval 1s;
        }
        mep 200 {
            interface ge-5/1/7.0;
            direction down;
            auto-discovery;
        }
    }
}
```

CFM on PE2

```
[edit interfaces]
ge-5/1/7 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}
ge-5/2/3 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}

[edit bridge-domains]
bridge-vlan2000 {
    domain-type bridge;
    interface ge-5/2/3.0;
    interface ge-5/1/7.0;
}

[edit protocols oam ethernet connectivity-fault-management]
maintenance-domain provider-outer {
    level 5;
    maintenance-association provider-outer-site1 {
        continuity-check {
            interval 1s;
        }
        mep 100 {
            interface ge-5/2/3.0;
            direction up;
            auto-discovery;
        }
    }
}
maintenance-domain provider-inner {
    level 3;
```



```
    maintenance-association provider-inner-site1 {
      continuity-check {
        interval 1s;
      }
      mep 100 {
        interface ge-5/1/7.0;
        direction down;
        auto-discovery;
      }
    }
  }
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Operations, Administration, and Maintenance on page 594
- Ethernet OAM Connectivity Fault Management on page 595
- Example: Configuring Ethernet CFM over VPLS on page 675
- Example: Configuring Ethernet CFM on Physical Interfaces on page 669

Example: Configuring Ethernet CFM over VPLS

In this example, both the customer and service provider are running Ethernet CFM over a VPLS and a multiprotocol label switching (MPLS) network. The network is shown in Figure 47 on page 676. The customer has configured Ethernet CFM on MX Series routers L2-CE1 and L2-CE2. The service provider has configured Ethernet CFM on MX Series routers PE1, P, and PE2.

NOTE: The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

The service provider is using CFM level 5 and the customer is using CFM level 7. The boundaries are marked with “up mep” and “down mep” CFM terminology in the figure.

Figure 47: Ethernet OAM with VPLS

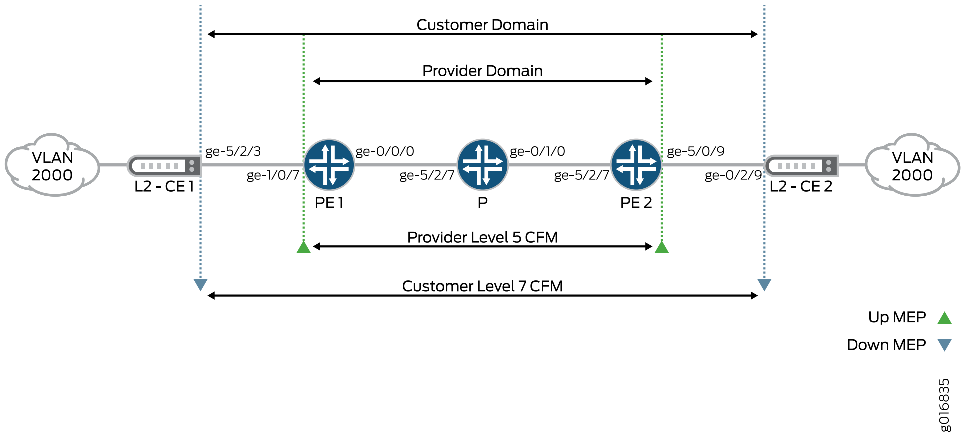

NOTE: The logical interfaces in a VPLS routing instance might have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. Normalization supports automatic mapping of VLANs and performs operations on VLAN tags to achieve the desired translation. See *Configuring a Normalized VLAN for Translation or Tagging*.

NOTE:

The following forwarding path considerations must be observed:

- Packet receive path:
  - This is the forwarding path for packets received on the interfaces.
  - 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.
- Packet transmit path:
  - Junos OS uses the router's hardware-based forwarding for CPU-generated packets.
  - For down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - In MX series routers, for up MEPs, the packets must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packets to be forwarded with this flood route.

The following are the configurations of the VPLS and CFM on the service provider routers.

Configuration of PE1

```
[edit chassis]
```

```
fpc 5 {
   pic 0 {
      tunnel-services {
         bandwidth 1g;
      }
   }
}

[edit interfaces]
ge-1/0/7 {
   encapsulation flexible-ethernet-services;
   vlan-tagging;
   unit 1 {
      encapsulation vlan-vpls;
      vlan-id 2000;
   }
}
ge-0/0/0 {
   unit 0 {
      family inet {
         address 10.200.1.1/24;
      }
      family mpls;
   }
}
lo0 {
   unit 0 {
      family inet {
         address 10.255.168.231/32 {
            primary;
         }
         address 127.0.0.1/32;
      }
   }
}

[edit routing-instances]
vpls-vlan2000 {
   instance-type vpls;
   vlan-id 2000;
   interface ge-1/0/7.1;
   route-distinguisher 10.255.168.231:2000;
   vrf-target target:1000:1;
   protocols {
      vpls {
         site-range 10;
         site vlan2000-PE1 {
            site-identifier 2;
         }
      }
   }
}

[edit protocols]
rsvp {
```

```
    interface ge-0/0/0.0;
}
mpls {
    label-switched-path PE1-to-PE2 {
        to 10.100.1.1;
    }
    interface ge-0/0/0.0;
}
bgp {
    group PE1-to-PE2 {
        type internal;
        local-address 10.200.1.1;
        family l2vpn {
            signaling;
        }
        local-as 65000;
        neighbor 10.100.1.1;
    }
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface ge-0/0/0.0;
    }
}
oam {
    ethernet {
        connectivity-fault-management {
            maintenance-domain customer-site1 {
                level 5;
                maintenance-association customer-site1 {
                    continuity-check {
                        interval 1s;
                    }
                    mep 100 {
                        interface ge-1/0/7.1;
                        direction up;
                        auto-discovery;
                    }
                }
            }
        }
    }
}
```

Configuration of PE2

```
[edit chassis]
fpc 5 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
```



```
    }
  }
}

[edit interfaces]
ge-5/0/9 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 2000;
  }
}
ge-5/2/7 {
  unit 0 {
    family inet {
      address 10.100.1.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.168.230/32 {
        primary;
      }
      address 127.0.0.1/32;
    }
  }
}

[edit routing-instances]
vpls-vlan2000 {
  instance-type vpls;
  vlan-id 2000;
  interface ge-5/0/9.1;
  route-distinguisher 10.255.168.230:2000;
  vrf-target target:1000:1;
  protocols {
    vpls {
      site-range 10;
      site vlan2000-PE2 {
        site-identifier 1;
      }
    }
  }
}

[edit protocols]
rsvp {
  interface ge-5/2/7.0;
}
mpls {
  label-switched-path PE2-to-PE1 {
```



```
        to 10.200.1.1;
      }
      interface ge-5/2/7.0;
    }
    bgp {
      group PE2-to-PE1 {
        type internal;
        local-address 10.100.1.1;
        family l2vpn {
          signaling;
        }
        local-as 65000;
        neighbor 10.200.1.1;
      }
    }
    ospf {
      traffic-engineering;
      reference-bandwidth 4g;
      area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
          disable;
        }
        interface ge-5/2/7.0;
      }
    }
    oam {
      ethernet {
        connectivity-fault-management {
          maintenance-domain customer-site1 {
            level 5;
            maintenance-association customer-site1 {
              continuity-check {
                interval 1s;
              }
              mep 200 {
                interface ge-5/0/9.1;
                direction up;
                auto-discovery;
              }
            }
          }
        }
      }
    }
```

Configuration of P router

MPLS only, no CFM needed:

```
[edit]
interfaces {
  ge-5/2/7 {
    # Connected to PE1
    unit 0 {
      family inet {
        address 10.200.1.10/24;
```



```
            }
            family mpls;
        }
    }
    ge-0/1/0 {
        # Connected to PE2
        unit 0 {
            family inet {
                address 10.100.1.10/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0{
            family inet {
                address 10.255.168.240/32;
            }
        }
    }
}

[edit]
protocols {
    rsvp {
        interface ge-0/1/0.0;
        interface ge-5/2/7.0;
    }
    mpls {
        interface ge-0/1/0.0;
        interface ge-5/2/7.0;
    }
    ospf {
        traffic-engineering;
        reference-bandwidth 4g;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
            interface ge-0/1/0.0;
            interface ge-5/2/7.0;
        }
    }
}
```

CFM on L2-CE1 Here is the configuration of CFM on L2-E1:

```
[edit interfaces]
ge-5/2/3 {
    vlan-tagging;
    unit 0 {
        vlan-id 2000;
    }
}
```

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 800 {
          interface ge-5/2/3.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
```

CFM on L2-CE2 Here is the configuration of CFM L2-CE2:

```
[edit interfaces]
ge-0/2/9 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 700 {
          interface ge-0/2/9.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Operations, Administration, and Maintenance on page 594
- Ethernet OAM Connectivity Fault Management on page 595
- Example: Configuring Ethernet CFM on Bridge Connections on page 671

- Example: Configuring Ethernet CFM on Physical Interfaces on page 669

CHAPTER 32

Configuring IEEE 802.3ah OAM Link-Fault Management

- Example: Configuring Ethernet LFM for CCC on page 715
- Example: Configuring Ethernet LFM for Aggregated Ethernet on page 716
- Example: Configuring Ethernet LFM with Loopback Support on page 718

IEEE 802.3ah OAM Link-Fault Management Overview

Ethernet interfaces capable of running at 100 Mbps or faster on EX Series switches, PTX Series, 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

Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured.

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.**

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured.

Related Documentation

- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Enabling Nonstop Routing for Ethernet Link Fault Management on Backup Routers on page 707
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

Understanding Ethernet OAM Link Fault Management for ACX Series Routers

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

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

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

The following OAM LFM features are supported on ACX Series routers:

- Discovery and Link Monitoring

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

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

- Remote Fault Detection

  Remote fault detection uses flags and events. Flags are used to convey the following:

  - **Link Fault** means a loss of signal
  - **Dying Gasp** means an unrecoverable condition such as a power failure. In this condition, the local peer informs the remote peer about the failure state. When the remote peer receives a dying-gasp PDU, it takes an action corresponding to the action profile configured with the **link-adjacency-loss** event.

    NOTE: ACX5096 and ACX5048 routers do not support dying-gasp.

    ACX Series routers can generate and receive dying-gasp packets. When LFM is configured on an interface, a dying-gasp PDU is generated for the interface on the following failure conditions:

    - Power failure
    - Packet Forwarding Engine panic or a crash

  - **Critical Event** means an unspecified vendor-specific critical event.

  You can specify the interval at which OAM PDUs are sent for fault detection.

NOTE: ACX Series routers support the receipt of dying-gasp packets, but cannot generate them.

- Remote Loopback Mode

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

  If a remote data terminal equipment (DTE) supports remote loopback mode, Junos OS can place the remote DTE into loopback mode. 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 and PAUSE frames are looped back. No changes are made to the frames. OAM PDUs continue to be sent and processed.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Configuring Ethernet Local Management Interface on page 627
- Ethernet OAM Connectivity Fault Management on page 595

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 1159
- action-profile on page 992

Configuring Ethernet 802.3ah OAM on PTX Series Packet Transport Routers

The IEEE 802.3ah standard for Operation, Administration, and Management (OAM) provides a specification for *Ethernet in the first mile (EFM)* connectivity. EFM defines how Ethernet can be transmitted over new media types using new Ethernet physical layer (PHY) interfaces. You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah OAM 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.

For Ethernet interfaces capable of running at 100 Mbps or faster, the IEEE 802.3ah OAM standard is supported on numerous Juniper Networks routers and switches. This topic describes configuration support for IEEE 802.3ah OAM features on PTX Series Packet Transport Routers.

Beginning in Junos OS Release 12.1, PTX Series routers support the following IEEE 802.3ah OAM features at the physical interface level:

- Discovery and link monitoring
- Fault signaling and detection
- Periodic packet management (PPM) processing

- Action profile support
- Graceful Routing Engine switchover (GRES)

To configure 802.3ah OAM support for Ethernet interfaces, include the **oam** statement at the **[edit protocols]** hierarchy level:

```
oam {
  ethernet {
    link-fault-management {
      interfaces {
        interface-name {
          pdu-interval interval;
          link-discovery (active | passive);
          pdu-threshold count;
        }
      }
    }
  }
}
```

Related Documentation

- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Configuring Link Discovery on page 692
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698

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 1223**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696

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 1220**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700

- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1298**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705

- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1299**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1159
- frame-error on page 1176
- frame-period on page 1177
- frame-period-summary on page 1178
- symbol-period on page 1383
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1271**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

Enabling Dying Gasp Functionality

Dying gasp means an unrecoverable condition such as a power failure. In this condition, the local peer informs the remote peer about the failure state. When the remote peer receives a dying-gasp PDU, it takes an action corresponding to the action profile configured with the **link-adjacency-loss** event. Dying gasp helps to avoid file system corruption.

NOTE: ACX5096 and ACX5048 routers do not support dying-gasp.

ACX Series routers can generate and receive dying-gasp packets. When LFM is configured on an interface, a dying-gasp PDU is generated for the interface on the following failure conditions:

- Power failure

- Packet Forwarding Engine panic or a crash

ACX Series routers support the following CLI statements to enable dying-gasp functionality:

- **dgasp-int**—Enables dying-gasp functionality.
- **dgasp-usb**—Resets USB port during dying-gasp event.

The **dgasp-int** and **dgasp-usb** CLI statements are added under the **[edit system]** hierarchy to enable dying-gasp functionality.

To enable dying-gasp functionality, you need to configure the **dgasp-int** and **dgasp-usb** CLI statements as shown below:

```
root@host% cli
root@host> configure
Entering configuration mode

[edit]
root@host# set system dgasp-int

[edit]
root@host# set system dgasp-usb

[edit]
root@host# commit

commit complete

[edit]
root@host# show system
dgasp-int;
dgasp-usb;
```

The dying-gasp functionality is disabled by default.

Related Documentation

- Understanding Ethernet OAM Link Fault Management for ACX Series Routers on page 687

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: Starting from Junos OS Release 14.2, 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.

Release History Table

Release	Description
14.2	Starting from Junos OS Release 14.2

Related Documentation

- action-profile on page 1068
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703

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 1067
- syslog on page 1384
- link-down on page 1222
- send-critical-event on page 1355

- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1220**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694

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 1325**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1222**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 710

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 1081**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694

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 1341**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693

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 1081
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693

Enabling Nonstop Routing for Ethernet Link Fault Management on Backup Routers

Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured. When the lfmd daemon runs on the backup Routing Engine as well, the link fault management states are kept in sync and so minimal effort is required by the lfmd daemon post switch over.

To enable Nonstop routing for Ethernet LFM on backup routers:

1. Enable graceful Routing Engine switchover. By default, GRES is disabled. To enable GRES, include the **graceful-switchover** statement at the [**edit chassis redundancy**] hierarchy level. By default, Nonstop routing is disabled. When you enable GRES, NSR is enabled.

   ```
   [edit chassis redundancy]
   user@host# set graceful-switchover
   ```

2. Synchronize the Routing Engine configuration. To synchronize the master Routing Engine configuration with the backup, include the **synchronize** statement at the [**edit system**] hierarchy level.

   ```
   [edit system]
   user@host# set commit synchronize
   ```

3. After enabling nonstop routing, commit the configuration.

   ```
   [edit routing options]
   user@host# commit
   ```


4. To verify if nonstop routing is enabled on the backup router, at the operational mode, use the **show oam ethernet link-fault-management** command on the master router and then the backup router. Because you have enabled synchronization, the output of the master router and the backup router is identical. However, the statistics maintained by the master router are not synchronized with the backup router..

   {master}
   user@host# show oam ethernet link-fault-management ge-0/2/0 detail

```
Interface: ge-0/2/0
    Status: Running, Discovery state: Send Any
    Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms
    Peer address: ac:4b:c8:81:90:a4
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM flags receive statistics:
      Critical event: 0, Dying gasp: 0, Link fault: 0
    OAM transmit statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 786, 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 received frame seconds 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
    Loopback tracking: Enabled, Loop status: Not Found
    Detect LOC: Enabled, LOC status: Not Found
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: unsupported, Link events: supported
      Variable requests: unsupported
Application profile statistics:
  Profile Name                Invoked     Executed
  LK_ADJ_LOSS100_1                     1        1
  LK_ADJ_LOSS100_2                     1        0
  LK_ADJ_LOSS100_3                     1        0
  LK_ADJ_LOSS101_1                     1        1
  LK_ADJ_LOSS101_2                     1        0
```



```
  LK_ADJ_LOSS101_3                        1        0
  LK_ADJ_LOSS106_1                        0        0
  LK_ADJ_LOSS106_2                        0        0
  LK_ADJ_LOSS106_3                        0        0
  LK_ADJ_LOSS107_1                        0        0
  LK_ADJ_LOSS107_2                        0        0
  LK_ADJ_LOSS107_3                        0        0
```

{backup}
user@host# show oam ethernet link-fault-management ge-0/2/0 detail

```
Interface: ge-0/2/0
    Status: Running, Discovery state: Send Any
    Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms
    Peer address: ac:4b:c8:81:90:a4
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM flags receive statistics:
      Critical event: 0, Dying gasp: 0, Link fault: 0
    OAM transmit statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 786, 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 received frame seconds 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
    Loopback tracking: Enabled, Loop status: Not Found
    Detect LOC: Enabled, LOC status: Not Found
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: unsupported, Link events: supported
      Variable requests: unsupported
Application profile statistics:
  Profile Name                 Invoked     Executed
  LK_ADJ_LOSS100_1                        0        0
  LK_ADJ_LOSS100_2                        0        0
  LK_ADJ_LOSS100_3                        0        0
  LK_ADJ_LOSS101_1                        0        0
```

```
LK_ADJ_LOSS101_2                              0          0
LK_ADJ_LOSS101_3                              0          0
LK_ADJ_LOSS106_1                              0          0
LK_ADJ_LOSS106_2                              0          0
LK_ADJ_LOSS106_3                              0          0
LK_ADJ_LOSS107_1                              0          0
LK_ADJ_LOSS107_2                              0          0
LK_ADJ_LOSS107_3                              0          0
```

NOTE: After the switchover, if issues are observed, use the **clear oam ethernet link-fault-management state** command for specific sessions. If the issue does not get resolved, restart the lfmd daemon.

Release History Table

Release	Description
17.3R1	Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured.

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- **show oam ethernet link-fault-management on page 2098**

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 1223**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691
- Configuring Link Discovery on page 692
- Configuring the OAM PDU Interval on page 693
- Configuring the OAM PDU Threshold on page 694
- Configuring Threshold Values for Local Fault Events on an Interface on page 694
- Disabling the Sending of Link Event TLVs on page 695
- Detecting Remote Faults on page 696
- Configuring an OAM Action Profile on page 698
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- Monitoring the Loss of Link Adjacency on page 701
- Monitoring Protocol Status on page 702
- Configuring Threshold Values for Fault Events in an Action Profile on page 703
- Applying an Action Profile on page 704
- Setting a Remote Interface into Loopback Mode on page 705
- Enabling Remote Loopback Support on the Local Interface on page 706

Example: Configuring IEEE 802.3ah OAM Support for an Interface on ACX Series

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

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

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2 or later for ACX Series routers.
- An ACX1000 or ACX2000 router.

Overview and Topology

In this example, you configure a 10-Gigabit Ethernet interface on an ACX Series router with 802.3ah OAM support, which includes: link discovery, protocol data units (PDUs), remote loopback, negotiation, and event thresholds.

Configuring IEEE 802.3ah OAM on an ACX Series Router

CLI Quick Configuration

To quickly configure IEEE 802.3ah Ethernet OAM, copy the following commands and paste them into the CLI:

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

Step-by-Step Procedure

To configure IEEE 802.3ah OAM support on an interface:

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

   ```
   [edit protocols oam ethernet link-fault-management]
   ```

   user@router1# **set** interface (OAM Link-Fault Management) **xe-0/0/0**

2. Specify that the interface initiates the discovery process by setting the link discovery mode to **active**:

   user@router# **set interface xe-0/0/0** link-discovery **active**

3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800:

   user@router# **set interface xe-0/0/0** pdu-interval **800**

4. Define the number of OAM PDUs to miss before an error is logged as 4:

   user@router# **set interface xe-0/0/0** pdu-threshold **4**

5. Configure the remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes:

   user@router# **set interface xe-0/0/0** remote-loopback

6. Configure remote loopback support for the local interface:

   user@router# **set interface xe-0/0/0** negotiation-options allow-remote-loopback

7. Set the threshold count for sending frame error events to 30:

   user@router# **set interface xe-0/0/0** event-thresholds frame-error **30**

8. Set the threshold count for sending frame period error events to 50:

   user@router# **set interface xe-0/0/0** event-thresholds frame-period **50**

9. Configure the threshold count for sending frame period summary error events to 40:

   user@router# **set interface xe-0/0/0** event-thresholds frame-period-summary **40**

10. Set the threshold count for sending symbol period events to 20:

    user@router# **set interface xe-0/0/0** event-thresholds symbol-period **20**

Results

Check the results of the configuration:

```
[edit]
user@router# show
```

```
[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 1223
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 689
- Enabling IEEE 802.3ah OAM Support on page 691

- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge

In this example, LFM is enabled on an IP link between the provider edge (PE) and customer edge (CE) interfaces. If the link goes down, the fault will be detected by LFM and the interfaces on both sides will be marked **Link-Layer-Down**. This results in notifications to various subsystems (for example, routing) which will take appropriate action.

The link running LFM is shown in Figure 48 on page 714.

Figure 48: Ethernet LFM Between Provider Edge and Customer Edge

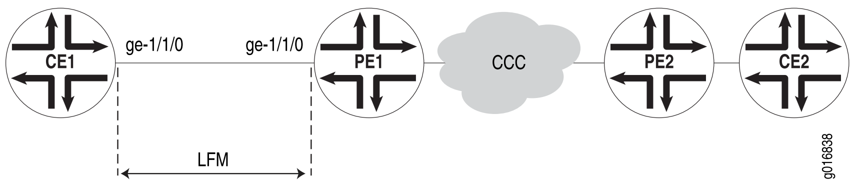

To configure Ethernet LFM on an IP link between PE and CE interfaces:

1. Configure LFM on the PE router:

```
[edit]
interfaces ge-1/1/0 {
   unit 0 {
      family inet {
         address 11.11.11.1/24;
      }
   }
}
protocols {
   oam {
      ethernet {
         link-fault-management {
            interface ge-1/1/0 {
               pdu-interval 1000;
               pdu-threshold 5;
            }
         }
      }
   }
}
```

2. Configure LFM on the CE router:

```
[edit]
interfaces ge-1/1/0 {
   unit 0 {
      family inet {
         address 11.11.11.2/24;
      }
   }
}
protocols {
   oam {
```



```
        ethernet {
            link-fault-management {
                interface ge-1/1/0 {
                    pdu-interval 1000;
                    pdu-threshold 5;
                }
            }
        }
    }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Example: Configuring Ethernet LFM for CCC on page 715
- Example: Configuring Ethernet LFM for Aggregated Ethernet on page 716
- Example: Configuring Ethernet LFM with Loopback Support on page 718

Example: Configuring Ethernet LFM for CCC

In this example, LFM is configured between two PEs (PE1 and PE2) connected using CCC. With LFM in place, a link fault will be detected immediately, instead of depending on routing protocols to find the fault on end-to-end CCC connection. This also helps in detecting the exact failed link instead of only finding that the end-to-end CCC connectivity has failed. Also, because LFM runs at the link-layer level, it does not need a IP address to operate and so can be used where bidirectional fault detection (BFD) cannot.

The links running LFM are shown in Figure 49 on page 715.

Figure 49: Ethernet LFM for CCC

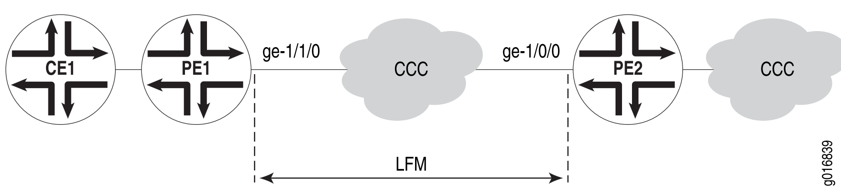

To configure Ethernet LFM between two PEs connected using CCC:

1. Configure LFM on the PE1 router with CCC:

```
[edit]
interfaces ge-1/1/0 {
    encapsulation ethernet-ccc;
    unit 0;
}
protocols {
    oam {
        ethernet {
            link-fault-management {
                interface ge-1/1/0 {
```

```
                pdu-interval 1000;
                pdu-threshold 5;
              }
            }
          }
        }
      }
```

2. Configure LFM on the PE2 router with CCC:

```
[edit]
interfaces ge-1/0/0 {
  encapsulation ethernet-ccc;
  unit 0;
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/0/0 {
          pdu-interval 1000;
          pdu-threshold 5;
        }
      }
    }
  }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 714
- Example: Configuring Ethernet LFM for Aggregated Ethernet on page 716
- Example: Configuring Ethernet LFM with Loopback Support on page 718

Example: Configuring Ethernet LFM for Aggregated Ethernet

In this example, LFM is configured on an aggregated Ethernet interface (AE0) between Router 1 and Router 2. When configured on aggregated Ethernet, LFM runs on all the individual member links. LFM is enabled or disabled on the member links as they are added or deleted from the aggregation group. The status of individual links is used to determine the status of the aggregated interface.

The use of LFM with aggregated Ethernet is shown in Figure 50 on page 717.

Figure 50: Ethernet LFM for Aggregated Ethernet

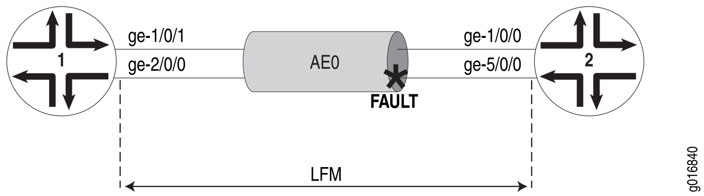

To configure LFM on an aggregated Ethernet interface between two routers:

1. Configure LFM on Router 1 for AE0:

```
[edit]
chassis {
   aggregated-devices {
      ethernet {
         device-count 1;
      }
   }
}
interfaces ge-1/0/1 {
   gigether-options {
      802.3ad ae0;
   }
}
interfaces ge-2/0/0 {
   gigether-options {
      802.3ad ae0;
   }
}
interfaces ae0 {
   unit 0 {
      family inet {
         address 11.11.11.2/24;
      }
   }
}
protocols {
   oam {
      ethernet {
         link-fault-management {
            interface ae0;
         }
      }
   }
}
```

2. Configure LFM on Router 2 for AE0:

```
[edit]
chassis {
   aggregated-devices {
      ethernet {
         device-count 1;
      }
```

```
    }
}
interfaces ge-1/0/0 {
    gigether-options {
        802.3ad ae0;
    }
}
interfaces ge-5/0/0 {
    gigether-options {
        802.3ad ae0;
    }
}
interfaces ae0 {
    unit 0 {
        family inet {
            address 11.11.11.1/24;
        }
    }
}
protocols {
    oam {
        ethernet {
            link-fault-management {
                interface ae0;
            }
        }
    }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 714
- Example: Configuring Ethernet LFM for CCC on page 715
- Example: Configuring Ethernet LFM with Loopback Support on page 718

Example: Configuring Ethernet LFM with Loopback Support

In this example, LFM is configured between provider edge (PE) router and the customer edge (CE) router. The PE router can put the CE router in remote loopback mode. This allows the PE to have all the traffic sent to the CE router looped back for diagnostics purposes, as shown in Figure 51 on page 719.

Figure 51: Ethernet LFM with Loopback Support

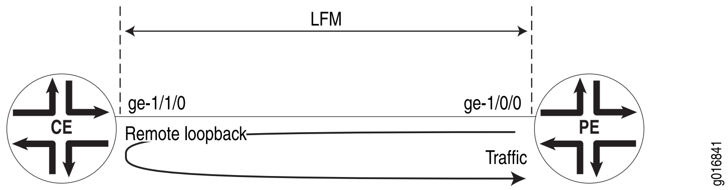

To configure LFM between a PE router and a CE router:

1. Configure LFM loopback on the PE router:

```
[edit]
interfaces ge-1/0/0 {
   unit 0 {
      family inet {
         address 11.11.11.1/24;
      }
   }
}
protocols {
   oam {
      ethernet {
         link-fault-management {
            interface ge-1/0/0 {
               pdu-interval 1000;
               pdu-threshold 5;
               remote-loopback;
            }
         }
      }
   }
}
```

2. Configure LFM loopback on the CE router:

```
[edit]
interfaces ge-1/1/0 {
   unit 0 {
      family inet {
         address 11.11.11.2/24;
      }
   }
}
protocols {
   oam {
      ethernet {
         link-fault-management {
            interface ge-1/1/0 {
               pdu-interval 1000;
               pdu-threshold 5;
               negotiation-options {
                  allow-remote-loopback;
               }
            }
```

```
            }
         }
      }
   }
```

NOTE: If the **negotiation options allow-remote-loopback** statement on the CE router is deleted before removing the CE router from remote loopback mode, traffic flow between the PE router and CE router is affected. Hence, delete the **remote-loopback** statement on the PE router before deleting the **negotiation-options allow-remote-loopback** statement on the CE router.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 686
- Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge on page 714
- Example: Configuring Ethernet LFM for CCC on page 715
- Example: Configuring Ethernet LFM for Aggregated Ethernet on page 716

CHAPTER 33

Configuring ITU-T Y.1731 Ethernet Service OAM

Ethernet Frame Delay Measurements Overview

ITU-T Y.1731 Frame Delay Measurement Feature

The IEEE 802.3-2005 standard for Ethernet Operations, Administration, and Maintenance (OAM) defines a set of link fault management mechanisms to detect and report link faults on a single point-to-point Ethernet LAN.

Junos OS supports key OAM standards that provide for automated end-to-end management and monitoring of Ethernet service by service providers:

- *IEEE Standard 802.1ag*, also known as “Connectivity Fault Management (CFM).”
- *ITU-T Recommendation Y.1731*, which uses different terminology than IEEE 802.1ag and defines Ethernet service OAM features for fault monitoring, diagnostics, and performance monitoring.

These capabilities allow operators to offer binding service-level agreements (SLAs) and generate new revenues from rate- and performance-guaranteed service packages that are tailored to the specific needs of their customers.

ACX Series routers support proactive and on-demand modes.

NOTE: **ACX5048 and ACX5096 routers supports only software-based time stamping for delay measurement.**

Ethernet CFM

The IEEE 802.1ag standard for connectivity fault management (CFM) defines mechanisms to provide for end-to-end Ethernet service assurance over any path, whether a single link or multiple links spanning networks composed of multiple LANs.

For Ethernet interfaces on M320, MX Series, and T Series routers, Junos OS supports the following key elements of the Ethernet CFM standard:

- Fault monitoring using the IEEE 802.1ag Ethernet OAM Continuity Check protocol
- Path discovery and fault verification using the IEEE 802.1ag Ethernet OAM Linktrace protocol
- Fault isolation using the IEEE 802.1ag Ethernet OAM Loopback protocol

In a CFM environment, network entities such as network operators, service providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management and end-to-end monitoring, and still avoid security breaches.

Figure 52 on page 724 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 52: Relationship of MEPs, MIPs, and Maintenance Domain Levels

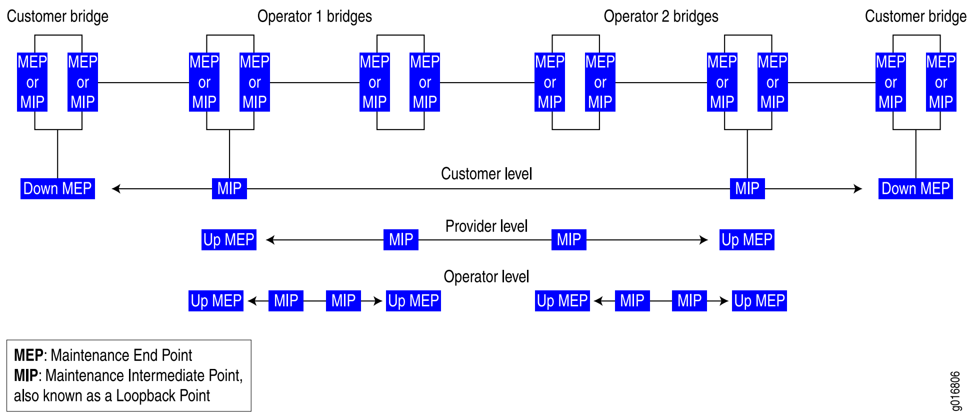

NOTE: On ACX Series routers, the maintenance intermediate points (MIP) is supported only on the ACX5048 and ACX5096 routers.

Ethernet Frame Delay Measurement

Two key objectives of OAM functionality are to measure quality-of-service attributes such as frame delay and frame delay variation (also known as “frame jitter”). Such measurements can enable you to identify network problems before customers are impacted by network defects.

Junos OS supports Ethernet frame delay measurement between MEPs configured on Ethernet physical or logical interfaces on MX Series routers. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor SLAs. Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. The Junos OS implementation of Ethernet frame delay measurement (ETH-DM) is fully compliant with the ITU-T Recommendation Y.1731, *OAM Functions and Mechanisms for Ethernet-based Networks*. The recommendation defines OAM mechanisms for operating and maintaining the network at the Ethernet service layer, which is called the "ETH layer" in ITU-T terminology.

MX Series routers with modular port concentrators (MPCs) and 10-Gigabit Ethernet MPCs with SFP+ support ITU-T Y.1731 functionality on VPLS for frame-delay and delay-variation.

NOTE: MX Series Virtual Chassis does not support Ethernet frame delay measurement (DM).

One-Way Ethernet Frame Delay Measurement

In one-way ETH-DM mode, a series of frame delay and frame delay variation values are calculated based on the time elapsed between the time a measurement frame is sent from the initiator MEP at one router and the time when the frame is received at the receiver MEP at the other router.

NOTE: ACX Series routers do not support one-way Ethernet frame delay measurement.

1DM Transmission

When you start a one-way frame delay measurement, the router sends 1DM frames—frames that carry the protocol data unit (PDU) for a one-way delay measurement—from the initiator MEP to the receiver MEP at the rate and for the number of frames you specify. The router marks each 1DM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

1DM Reception

When an MEP receives a 1DM frame, the router that contains the receiver MEP measures the one-way delay for that frame (the difference between the time the frame was received and the timestamp contained in the frame itself) and the delay variation (the difference between the current and previous delay values).

One-Way ETH-DM Statistics

The router that contains the receiver MEP stores each set of one-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

One-Way ETH-DM Frame Counts

Each router counts the number of one-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number of 1DM frames sent.
- For a receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

Synchronization of System Clocks

The accuracy of one-way delay calculations depends on close synchronization of the system clocks at the initiator MEP and receiver MEP.

The accuracy of one-way delay variation is not dependent on system clock synchronization. Because delay variation is simply the difference between consecutive one-way delay values, the out-of-phase period is eliminated from the frame jitter values.

NOTE: **For a given one-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only on the router that contains the receiver MEP.**

Two-Way Ethernet Frame Delay Measurement

In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

DMM Transmission

When you start a two-way frame delay measurement, the router sends delay measurement message (DMM) frames— frames that carry the PDU for a two-way ETH-DM request—from the initiator MEP to the responder MEP at the rate and for the number of frames you specify. The router marks each DMM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

DMR Transmission

When an MEP receives a DMM frame, the responder MEP responds with a delay measurement reply (DMR) frame, which carries ETH-DM reply information and a copy of the timestamp contained in the DMM frame.

DMR Reception

When an MEP receives a valid DMR, the router that contains the MEP measures the two-way delay for that frame based on the following sequence of timestamps:

1. TI_{TxDMM}
2. TR_{RxDMM}
3. TR_{TxDMR}
4. TI_{RxDMR}

A two-way frame delay is calculated as follows:

$$[TI_{RxDMR} - TI_{TxDMM}] - [TR_{TxDMR} - TR_{RxDMM}]$$

The calculation show that frame delay is the difference between the time at which the initiator MEP sends a DMM frame and the time at which the initiator MEP receives the associated DMR frame from the responder MEP, minus the time elapsed at the responder MEP.

The delay variation is the difference between the current and previous delay values.

Two-Way ETH-DM Statistics

The router that contains the initiator MEP stores each set of two-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

Two-Way ETH-DM Frame Counts

Each router counts the number of two-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number DMM frames transmitted, the number of valid DMR frames received, and the number of invalid DMR frames received.
- For a responder MEP, the router counts the number of DMR frames sent.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

NOTE: **For a given two-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only at the router that contains the initiator MEP.**

Choosing Between One-Way and Two-Way ETH-DM

One-way frame delay measurement requires that the system clocks at the initiator MEP and receiver MEP are closely synchronized. Two-way frame delay measurement does not require synchronization of the two systems. If it is not practical for the clocks to be synchronized, two-way frame delay measurements are more accurate.

When two systems are physically close to each other, their one-way delay values are very high compared to their two-way delay values. One-way delay measurement requires that the timing for the two systems be synchronized at a very granular level, and MX Series routers currently do not support this granular synchronization.

Restrictions for Ethernet Frame Delay Measurement

The following restrictions apply to the Ethernet frame delay measurement feature:

- The ETH-DM feature is not supported on label-switched interface (LSI) pseudowires.

  The ETH-DM feature is supported on aggregated Ethernet interfaces.

- Hardware-assisted timestamping for ETH-DM frames in the reception path is only supported for MEP interfaces on Enhanced DPCs and Enhanced Queuing DPCs in MX Series routers. For information about hardware-assisted timestamping, see "Guidelines for Configuring Routers to Support an ETH-DM Session" on page 768 and "Enabling the Hardware-Assisted Timestamping Option" on page 779.
- Ethernet frame delay measurements can be triggered only when the distributed periodic packet management daemon (**ppm**) is enabled. For more information about this limitation, see "Guidelines for Configuring Routers to Support an ETH-DM Session" on page 768 and "Ensuring That Distributed ppm Is Not Disabled" on page 776.
- You can monitor only one session at a time to the same remote MEP or MAC address. For more information about starting an ETH-DM session, see "Starting an ETH-DM Session" on page 780.
- ETH-DM statistics are collected at only one of the two peer routers in the ETH-DM session. For a one-way ETH-DM session, you can display frame ETH-DM statistics at the receiver MEP only, using ETH-DM-specific **show** commands. For a two-way ETH-DM session, you can display frame delay statistics at the initiator MEP only, using the same ETH-DM-specific **show** commands. For more information, see "Managing ETH-DM Statistics and ETH-DM Frame Counts" on page 784.
- ETH-DM frame counts are collected at both MEPs and are stored in the respective CFM databases.
- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. Therefore, the collection of ETH-DM statistics and ETH-DM frame counters has to be restarted, after the switchover is complete. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *Junos OS High Availability Library for Routing Devices*.
- Accuracy of frame delay statistics is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Related Documentation

- Ethernet Frame Loss Measurement Overview on page 729
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768
- Guidelines for Starting an ETH-DM Session on page 769
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 771
- On-Demand Mode for SLA Measurement on page 732
- Proactive Mode for SLA Measurement on page 733

Ethernet Frame Loss Measurement Overview

The key objectives of the OAM functionality are to measure quality-of-service attributes such as frame delay, frame delay variation (also known as "frame jitter"), and frame loss. Such measurements enable you to identify network problems before customers are impacted by network defects. For more information about Ethernet frame delay measurement, see "Ethernet Frame Delay Measurements Overview" on page 723.

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on MX Series routers and is presently supported only for VPWS service. ETH-LM is used by operators to collect counter values applicable for ingress and egress service frames. These counters maintain a count of transmitted and received data frames between a pair of MEPs. Ethernet frame loss measurement is performed by sending frames with ETH-LM information to a peer MEP and similarly receiving frames with ETH-LM information from the peer MEP. This type of frame loss measurement is also known as single-ended Ethernet loss measurement.

NOTE: MX Series Virtual Chassis does not support Ethernet frame loss measurement (ETH-LM).

ETH-LM supports the following frame loss measurements:

- Near-end frame loss measurement—Measurement of frame loss associated with ingress data frames.
- Far-end frame loss measurement—Measurement of frame loss associated with egress data frames.

NOTE: The proactive and dual-ended loss measurement functionality of ITU-T Y1731 is not supported on the ACX Series routers.

The ETH-LM feature is supported on aggregated Ethernet interfaces.

NOTE: Starting Junos OS Release 16.1, the Ethernet loss measurement (ETH-LM) results are inaccurate when connectivity fault management (CFM) and performance monitoring (PM) PDUs received locally at a maintenance endpoint (MEP) as classified as belonging to the yellow class or a packet loss priority (PLP) of medium-high. This problem of incorrect results is specific to Ethernet loss measurement for CFM sessions of down MEPs. The Ethernet loss measurement statistics are inaccurate in the following scenarios:

- Ethernet loss measurement is working on a CFM session for a MEP in down state
- CFM PDUs received on the logical interface of the down MEP are classified by the classifier as yellow or medium-high PLP
- A packet is identified as yellow when the input classifier marks the PLP as medium-high.

The problem of discrepancies with Ethernet loss measurement results is not observed when you configure Ethernet loss measurement in colorless mode. To avoid this problem of inaccurate loss measurement results, provision all local CFM PDUs as green or with the PLP as high.

NOTE: Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the **performance-monitoring** statement and its substatements at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level) is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

Release History Table

Release	Description
16.1	Starting Junos OS Release 16.1, the Ethernet loss measurement (ETH-LM) results are inaccurate when connectivity fault management (CFM) and performance monitoring (PM) PDUs received locally at a maintenance endpoint (MEP) as classified as belonging to the yellow class or a packet loss priority (PLP) of medium-high.
16.1	Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the **performance-monitoring** statement and its substatements at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level) is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

Related Documentation

- Managing Continuity Measurement Statistics on page 803
- On-Demand Mode for SLA Measurement on page 732
- Proactive Mode for SLA Measurement on page 733

- Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs on page 828
- Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs on page 840

Service-Level Agreement Measurement

Service-level agreement (SLA) measurement is the process of monitoring the bandwidth, delay, delay variation (jitter), continuity, and availability of a service (E-Line or E-LAN). It enables you to identify network problems before customers are impacted by network defects.

NOTE:
The Ethernet VPN services can be classified into:

- Peer-to-peer-services (E-Line services)—The E-Line services are offered using MPLS-based Layer 2 VPN virtual private wire service (VPWS).
- Multipoint-to-multipoint services (E-LAN services)—The E-LAN services are offered using MPLS-based virtual private LAN service (VPLS).

For more information, see the *Junos VPNs Configuration Guide*.

In Junos OS, SLA measurements are classified into:

- On-demand mode—In on-demand mode, the measurements are triggered through the CLI. For more information, see "On-Demand Mode for SLA Measurement" on page 732.
- Proactive mode—In proactive mode, the measurements are triggered by an iterator application. For more information, see "Proactive Mode for SLA Measurement" on page 733.

For more information about frame delay measurement, see "Ethernet Frame Delay Measurements Overview" on page 723. For more information about frame loss measurement, see "Ethernet Frame Loss Measurement Overview" on page 729. Note that Ethernet frame delay measurement and Ethernet frame loss measurement are not supported on the **ae** interface.

Related Documentation

- Proactive Mode for SLA Measurement on page 733.
- On-Demand Mode for SLA Measurement on page 732.

On-Demand Mode for SLA Measurement

In on-demand mode, the measurements are triggered by the user through the CLI.

When the user triggers the delay measurement through the CLI, the delay measurement request that is generated is as per the frame formats specified by the ITU-T Y.1731 standard. For two-way delay measurement, the server-side processing can be delegated to the Packet Forwarding Engine to prevent overloading on the Routing Engine. For more information, see "Configuring Routers to Support an ETH-DM Session" on page 775. When the server-side processing is delegated to the Packet Forwarding Engine, the delay measurement message (DMM) frame **receive** counters and delay measurement reply (DMR) frame **transmit** counters are not displayed by the **show** command.

When the user triggers the loss measurement through the CLI, the router sends the packets in standard format along with the loss measurement TLV. By default, the **session-id-tlv** argument is included in the packet to allow concurrent loss measurement sessions from same local MEP. You can also disable the session ID TLV by using the **no-session-id-tlv** argument.

Single-ended ETH-LM is used for on-demand operation, administration, and maintenance purposes. An MEP sends frames with ETH-LM request information to its peer MEP and receives frames with ETH-LM reply information from its peer MEP to carry out loss measurements. The protocol data unit (PDU) used for a single-ended ETH-LM request is referred to as a loss measurement message (LMM) and the PDU used for a single-ended ETH-LM reply is referred to as a loss measurement reply (LMR).

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Ethernet Frame Loss Measurement Overview on page 729
- Proactive Mode for SLA Measurement on page 733
- Configuring Routers to Support an ETH-DM Session on page 775.

Proactive Mode for SLA Measurement

In proactive mode, SLA measurements are triggered by an iterator application. An iterator is designed to periodically transmit SLA measurement packets in form of ITU-Y.1731-compliant frames for two-way delay measurement or loss measurement on MX Series routers. This mode differs from on-demand SLA measurement, which is user initiated. The iterator sends periodic delay or loss measurement request packets for each of the connections registered to it. Iterators make sure that measurement cycles do not occur at the same time for the same connection to avoid CPU overload. Junos OS supports proactive mode for VPWS. For an iterator to form a remote adjacency and to become functionally operational, the continuity check message (CCM) must be active between the local and remote MEP configurations of the connectivity fault management (CFM). Any change in the iterator adjacency parameters resets the existing iterator statistics and restarts the iterator. Here, the term adjacency refers to a pairing of two endpoints (either connected directly or virtually) with relevant information for mutual understanding, which is used for subsequent processing. For example, the iterator adjacency refers to the iterator association between the two endpoints of the MEPs.

For every DPC or MPC, only 30 iterator instances for a cycle time value of 10 milliseconds (ms) are supported. In Junos OS, 255 iterator profile configurations and 2000 remote MEP associations are supported.

Iterators with cycle time value less than 100 ms are supported only for infinite iterators, whereas the iterators with cycle time value greater than 100 ms are supported for both finite and infinite iterators. Infinite iterators are iterators that run infinitely until the iterator is disabled or deactivated manually.

NOTE: **ACX5048 and ACX5096 routers supports iterator cycle time of only 1 second and above.**

A VPWS service configured on a router is monitored for SLA measurements by registering the connection (here, the connection is a pair of remote and local MEPs) on an iterator and then initiating periodic SLA measurement frame transmission on those connections. The end-to-end service is identified through a maintenance association end point (MEP) configured at both ends.

For two-way delay measurement and loss measurement, an iterator sends a request message for the connection in the list (if any) and then sends a request message for the connection that was polled in the former iteration cycle. The back-to-back request messages for the SLA measurement frames and their responses help in computing delay variation and loss measurement.

The Y.1731 frame transmission for a service attached to an iterator continues endlessly unless intervened and stopped by an operator or until the iteration-count condition is met. To stop the iterator from sending out any more proactive SLA measurement frames, the operator must perform one of the following tasks:

- Enable the **deactivate sla-iterator-profile** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance association *ma-name* mep *mep-id* remote-mep *mep-id*]** hierarchy level. For more information, see "Verifying the Configuration of an Iterator Profile" on page 756.
- Provision a **disable** statement under the corresponding iterator profile at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]** hierarchy level. For more information, see "Configuring an Iterator Profile" on page 754.

Ethernet Delay Measurements and Loss Measurement by Proactive Mode

In two-way delay measurement, the delay measurement message (DMM) frame is triggered through an iterator application. The DMM frame carries an iterator type, length, and value (TLV) in addition to the fields described in standard frame format and the server copies the iterator TLV from the DMM frame to the delay measurement reply (DMR) frame.

In one-way delay variation computation using the two-way delay measurement method, the delay variation computation is based on the timestamps that are present in the DMR frame (and not the 1DM frame). Therefore, there is no need for client-side and server-side clocks to be in sync. Assuming that the difference in their clocks remains constant, the one-way delay variation results are expected to be fairly accurate. This method also eliminates the need to send separate 1DM frames just for the one-way delay variation measurement purpose.

In proactive mode for loss measurement, the router sends packets in standard format along with loss measurement TLV and iterator TLV.

Related Documentation

- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Ethernet Frame Delay Measurements Overview on page 723
- Ethernet Frame Loss Measurement Overview on page 729
- Verifying the Configuration of an Iterator Profile on page 756
- Managing Iterator Statistics on page 759
- On-Demand Mode for SLA Measurement on page 732

Ethernet Failure Notification Protocol Overview

The Failure Notification Protocol (FNP) is a failure notification mechanism that detects failures in Point-to-Point Ethernet transport networks on MX Series routers. If a node link fails, FNP detects the failure and sends out FNP messages to the adjacent nodes that a circuit is down. Upon receiving the FNP message, nodes can redirect traffic to the protection circuit.

NOTE: FNP is supported on E-Line services only.

An E-Line service provides a secure Point-to-Point Ethernet connectivity between two user network interfaces (UNIs). E-Line services are a protected service and each service has a working circuit and protection circuit. CFM is used to monitor the working and protect paths. CCM intervals result in failover time in hundreds of milliseconds or a few seconds. FNP provides service circuit failure detection and propagation in less than 50ms and provide 50ms failover for E-Line services.

The MX router acts as a PE node and handles the FNP messages received on the management VLAN and the FNP messages received on both the Ethernet interfaces and PWs created for the management VPLS. MX-series routers do not initiate FNP messages and responds only to FNP messages generated by devices in the Ethernet Access network. FNP can be enabled only on logical interfaces that are part of a VPLS routing instance, and no physical interfaces in that VPLS routing instance should have CCM configured. FNP can be enabled only on one logical interface per physical interface.

All E-Line services are configured as layer 2 circuits with edge protection. A VLAN associated with the working circuit or protection circuit must map to a logical interface. No trunk port or access port is supported in the ring link for VLANs used by E-LINE services. FNP does not control the logical interface associated with protection circuit. Only E-Line service whose termination point is not in an MX node is controlled by FNP.

FNP supports graceful restart and the Graceful Routing Engine switchover (GRES) features.

Related Documentation

- Configuring the Failure Notification Protocol on page 804
- **show oam ethernet fnp interface on page 2092**
- **show oam ethernet fnp status on page 2096**
- **show oam ethernet fnp messages on page 2094**
- **connectivity-fault-management on page 1113**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

Ethernet Synthetic Loss Measurement Overview

Ethernet synthetic loss measurement (ETH-SLM) is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic. Each maintenance association end point (MEP) performs frame loss measurements, which contribute to unavailable time.

A near-end frame loss specifies frame loss associated with ingress data frames and a far-end frame loss specifies frame loss associated with egress data frames. Both near-end and far-end frame loss measurements contribute to near-end severely errored seconds and far-end severely errored seconds that are used in combination to determine unavailable time. ETH-SLM is performed using synthetic loss message (SLM) and synthetic loss reply (SLR) frames. ETH-SLM facilitates each MEP to perform near-end and far-end synthetic frame loss measurements by using synthetic frames because a bidirectional service is defined as unavailable if either of the two directions is determined to be unavailable.

There are the two types of frame loss measurement, defined by the ITU-T Y.1731 standards, ETH-LM and ETH-SLM. Junos OS supports only single-ended ETH-SLM. In single-ended ETH-SLM, each MEP sends frames with the ETH-SLM request information to its peer MEP and receives frames with ETH-SLM reply information from its peer MEP to perform synthetic loss measurements. Single-ended ETH-SLM is used for proactive or on-demand OAM to perform synthetic loss measurements applicable to point-to-point Ethernet connection. This method allows a MEP to initiate and report far-end and near-end loss measurements associated with a pair of MEPs that are part of the same maintenance entity group (MEG).

NOTE: **MX Series Virtual Chassis does not support Ethernet synthetic loss measurement (ETH-SLM).**

Single-ended ETH-SLM is used to perform on-demand or proactive tests by initiating a finite amount of ETH-SLM frames to one or multiple MEP peers and receiving the ETH-SLM reply from the peers. The ETH-SLM frames contain the ETH-SLM information that is used to measure and report both near-end and far-end synthetic loss measurements. Service-level agreement (SLA) measurement is the process of monitoring the bandwidth, delay, delay variation (jitter), continuity, and availability of a service. It enables you to identify network problems before customers are impacted by network defects. In proactive mode, SLA measurements are triggered by an iterator application. An iterator is designed to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for synthetic frame loss measurement . This mode differs from on-demand SLA measurement, which is user initiated. In on-demand mode, the measurements are triggered by the user through the CLI. When the user triggers the ETH-SLM through the CLI, the SLM request that is generated is as per the frame formats specified by the ITU-T Y.1731 standard.

NOTE: **ACX5048 and ACX5096 routers support ETH-SLM for Layer 2 services.**

Related Documentation

- Transmission of ETH-SLM Messages on page 740
- Format of ETH-SLM Messages on page 738
- Guidelines for Configuring ETH-SLM on page 742
- Scenarios for Configuration of ETH-SLM on page 737

- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Starting an On-Demand ETH-SLM Session on page 747
- Troubleshooting Failures with ETH-SLM on page 752

Scenarios for Configuration of ETH-SLM

ETH-SLM measures near-end and far-end frame loss between two MEPs that are part of the same MEG level. You can configure ETH-SLM to measure synthetic loss for both upward-facing or upstream MEP and downward-facing or downstream MEP. This section describes the following scenarios for the operation of ETH-SLM:

Upstream MEP in MPLS Tunnels

Consider a scenario in which a MEP is configured between the user network interfaces (UNIs) of two MX Series routers, MX1 and MX2, in the upstream direction. MX1 and MX2 are connected over an MPLS core network. ETH-SLM measurements are performed between the upstream MEP in the path linking the two routers. Both MX1 and MX2 can initiate on-demand or proactive ETH-SLM, which can measure both far-end and near-end loss at MX1 and MX2, respectively. The two UNIs are connected using MPLS-based Layer 2 VPN virtual private wire service (VPWS).

Downstream MEP in Ethernet Networks

Consider a scenario in which a MEP is configured between two MX Series routers, MX1 and MX2, on the Ethernet interfaces in the downstream direction. MX1 and MX2 are connected in an Ethernet topology and downstream MEP is configured toward the Ethernet network. ETH-SLM measurements are performed between the downstream MEP in the path linking the two routers. ETH-SLM can be measured in the path between these two routers.

Consider another scenario in which a MEP is configured in the downstream direction and service protection for a VPWS over MPLS is enabled by specifying a working path or protect path on the MEP. Service protection provides end-to-end connection protection of the working path in the event of a failure. To configure service protection, you must create two separate transport paths—a working path and a protect path. You can specify the working path and protect path by creating two maintenance associations. To associate the maintenance association with a path, you must configure the MEP interface in the maintenance association and specify the path as working or protect.

In a sample topology, an MX Series router, MX1, is connected to two other MX Series routers, MX2 and MX3, over an MPLS core. The connectivity fault management (CFM) session between MX1 and MX2 is the working path on the MEP and the CFM session between MX1 and MX3 is the protect path on the MEP. MX2 and MX3 are, in turn, connected on Ethernet interfaces to MX4 in the access network. Downstream MEP is configured between MX1 and MX4 that passes through MX2 (working CFM session) and also between MX1 and MX4 that passes through MX3 (protected CFM session). ETH-SLM is performed between these downstream MEPs. In both the downstream MEPs, the configuration is performed on MX1 and MX4 UNIs, similar to upstream MEP.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Transmission of ETH-SLM Messages on page 740
- Format of ETH-SLM Messages on page 738
- Guidelines for Configuring ETH-SLM on page 742
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Starting an On-Demand ETH-SLM Session on page 747
- Troubleshooting Failures with ETH-SLM on page 752

Format of ETH-SLM Messages

Synthetic loss messages (SLMs) support single-ended Ethernet synthetic loss measurement (ETH-SLM) requests. This topic contains the following sections that describe the formats of the SLM protocol data units (PDUs), SLR PDUs, and the data iterator type length value (TLV).

SLM PDU Format

The SLM PDU format is used by a MEP to transmit SLM information. The following components are contained in SLM PDUs:

- Source MEP ID—Source MEP ID is a 2-octet field where the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.
- Test ID—Test ID is a 4-octet field set by the transmitting MEP and is used to identify a test when multiple tests run simultaneously between MEPs (including both concurrent on-demand and proactive tests).
- TxFCf—TxFCf is a 4-octet field that carries the number of SLM frames transmitted by the MEP toward its peer MEP.

The following are the fields in an SLM PDU:

- MEG Level—Configured maintenance domain level in the range 0–7.
- Version—0.
- OpCode—Identifies an OAM PDU type. For SLM, it is 55.
- Flags—Set to all zeros.
- TLV Offset—16.
- Source MEP ID—A 2-octet field used to identify the MEP transmitting the SLM frame. In this 2-octet field, the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.

- RESV—Reserved fields are set to all zeros.
- Test ID—A 4-octet field set by the transmitting MEP and used to identify a test when multiple tests run simultaneously between MEPs (including both concurrent on-demand and proactive tests).
- TxFCf—A 4-octet field that carries the number of SLM frames transmitted by the MEP toward its peer MEP.
- Optional TLV—A data TLV may be included in any SLM transmitted. For the purpose of ETH-SLM, the value part of data TLV is unspecified.
- End TLV—All zeros octet value.

SLR PDU Format

The synthetic loss reply (SLR) PDU format is used by a MEP to transmit SLR information. The following are the fields in an SLR PDU:

- MEG Level—A 3-bit field the value of which is copied from the last received SLM PDU.
- Version—A 5-bit field the value of which is copied from the last received SLM PDU.
- OpCode—Identifies an OAM PDU type. For SLR, it is set as 54.
- Flags—A 1-octet field copied from the SLM PDU.
- TLV Offset—A 1-octet field copied from the SLM PDU.
- Source MEP ID—A 2-octet field copied from the SLM PDU.
- Responder MEP ID—A 2-octet field used to identify the MEP transmitting the SLR frame.
- Test ID—A 4-octet field copied from the SLM PDU.
- TxFCf—A 4-octet field copied from the SLM PDU.
- TxFCb—A 4 octet field. This value represents the number of SLR frames transmitted for this test ID.
- Optional TLV—The value is copied from the SLM PDU, if present.
- End TLV—A 1-octet field copied from the SLM PDU.

Data Iterator TLV Format

The data iterator TLV specifies the data TLV portion of the Y.1731 data frame. The MEP uses a data TLV when the MEP is configured to measure delay and delay variation for different frame sizes. The following are the fields in a data TLV:

- Type—Identifies the TLV type; value for this TLV type is Data (3).
- Length—Identifies the size, in octets, of the Value field containing the data pattern. The maximum value of the Length field is 1440.
- Data pattern—An *n*-octet (*n* denotes length) arbitrary bit pattern. The receiver ignores it.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Transmission of ETH-SLM Messages on page 740
- Guidelines for Configuring ETH-SLM on page 742
- Scenarios for Configuration of ETH-SLM on page 737
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Starting an On-Demand ETH-SLM Session on page 747
- Troubleshooting Failures with ETH-SLM on page 752

Transmission of ETH-SLM Messages

The ETH-SLM functionality can process multiple synthetic loss message (SLM) requests simultaneously between a pair of MEPs. The session can be a proactive or an on-demand SLM session. Each SLM request is identified uniquely by a test ID.

A MEP can send SLM requests or respond to SLM requests. A response to an SLM request is called a synthetic loss reply (SLR). After a MEP determines an SLM request by using the test ID, the MEP calculates the far-end and near-end frame loss on the basis of the information in the SLM message or the SLM protocol data unit (PDU).

A MEP maintains the following local counters for each test ID and for each peer MEP being monitored in a maintenance entity for which loss measurements are to be performed:

- TxFCl—Number of synthetic frames transmitted toward the peer MEP for a test ID. A source MEP increments this number for successive transmission of synthetic frames with ETH-SLM request information while a destination or receiving MEP increments this value for successive transmission of synthetic frames with the SLR information.
- RxFCl—Number of synthetic frames received from the peer MEP for a test ID. A source MEP increments this number for successive reception of synthetic frames with SLR information while a destination or receiving MEP increments it for successive reception of synthetic frames with ETH-SLM request information.

The following sections describe the phases of processing of SLM PDUs to determine synthetic frame loss:

Initiation and Transmission of SLM Requests

A MEP periodically transmits an SLM request with the OpCode field set as 55. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TxFCl is sent in the packet.

No synchronization is required of the test ID value between initiating and responding MEPs because the test ID is configured at the initiating MEP, and the responding MEP uses the test ID it receives from the initiating MEP. Because ETH-SLM is a sampling technique, it is less precise than counting the service frames. Also, the accuracy of measurement depends on the number of SLM frames used or the period for transmitting SLM frames.

Reception of SLMs and Transmission of SLRs

After the destination MEP receives a valid SLM frame from the source MEP, an SLR frame is generated and transmitted to the requesting or source MEP. The SLR frame is valid if the MEG level and the destination MAC address match the receiving MEP's MAC address. All the fields in the SLM PDUs are copied from the SLM request except for the following fields:

- The source MAC address is copied to the destination MAC address and the source address contains the MEP's MAC address.
- The value of the OpCode field is changed from SLM to SLR (54).
- The responder MEP ID is populated with the MEP's MEP ID.
- TxFCb is saved with the value of the local counter RxFCl at the time of SLR frame transmission.
- An SLR frame is generated every time an SLM frame is received; therefore, RxFCl in the responder is equal to the number of SLM frames received and also equal to the number of SLR frames sent. At the responder or receiving MEP, RxFCl equals TxFCl.

Reception of SLRs

After an SLM frame (with a given TxFCf value) is transmitted, a MEP expects to receive a corresponding SLR frame (carrying the same TxTCf value) within the timeout value from its peer MEP. SLR frames that are received after the timeout value (5 seconds) are discarded. With the information contained in SLR frames, a MEP determines the frame loss for the specified measurement period. The measurement period is a time interval during which the number of SLM frames transmitted is statistically adequate to make a measurement at a given accuracy. A MEP uses the following values to determine near-end and far-end frame loss during the measurement period:

- Last received SLR frame's TxFCf and TxFCb values and the local counter RxFCl value at the end of the measurement period. These values are represented as TxFCf[tc], TxFCb[tc], and RxFCl[tc], where tc is the end time of the measurement period.
- SLR frame's TxFCf and TxFCb values of the first received SLR frame after the test starts and local counter RxFCl at the beginning of the measurement period. These values are represented as TxFCf[tp], TxFCb[tp], and RxFCl[tp], where tp is the start time of the measurement period.

For each SLR packet that is received, the local RxFCl counter is incremented at the sending or source MEP.

Computation of Frame Loss

Synthetic frame loss is calculated at the end of the measurement period on the basis of the value of the local counters and the information from the last frame received. The last received frames contains the TxFCf and TxFCb values. The local counter contains the RxFCl value. Using these values, frame loss is determined using the following formula:

Frame loss (far-end) = TxFCf – TxFCb

Frame loss (near-end) = TxFCb – RxFCl

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Format of ETH-SLM Messages on page 738
- Guidelines for Configuring ETH-SLM on page 742
- Scenarios for Configuration of ETH-SLM on page 737
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Starting an On-Demand ETH-SLM Session on page 747
- Troubleshooting Failures with ETH-SLM on page 752

Guidelines for Configuring ETH-SLM

Keep the following points in mind when you configure the ETH-SLM functionality:

- The monitoring application for Ethernet OAM is initiated in the master Routing Engine. When a stateful switchover process occurs, the monitoring application is disabled. For on-demand ETH-SLM, graceful Routing Engine switchover (GRES) support is not applicable. For proactive ETH-SLM, the service-level agreement (SLA) iterators are restored during a stateful switchover process. If the adjacencies do not time out, the ETH-SLM statistics are preserved and proactive ETH-SLM supports GRES.
- ETH-SLM is initiated only when the MEP session is up. Unified in-service software upgrade (ISSU) support for ETH-SLM depends on the unified ISSU support for CFM. For CFM, unified ISSU is supported using the loss threshold TLV to avoid CFM connectivity loss during the upgrade. The receiving or the destination MEP increases the threshold time during the termination of sessions. If you start a unified ISSU operation when on-demand ETH-SLM is in progress, the SLM request and reply messages are lost at the local Packet Forwarding Engine.

  When an on-demand ETH-SLM is requested, if the local source MEP undergoes a unified ISSU, a message is displayed stating that the MEP is undergoing a unified ISSU. If the remote MEP is undergoing a unified ISSU (detected through the loss threshold TLV), a message is displayed stating that the remote MEP is undergoing a unified ISSU. Also, if it is not possible to identify whether unified ISSU is in progress on a remote MEP, the SLM packets are lost at the system where unified ISSU is in progress and the loss calculation results do not provide a valid cause for the loss. Unified ISSU is not supported for both on-demand and proactive ETH-SLM.
- The maximum number of SLA iterator profiles that can be configured in the system is 255.

- ETH-SLM is not supported for virtual private LAN service (VPLS) (point-to-multipoint measurements are not supported). The ETH-SLM frames are not generated with multicast class 1 destination address. Similarly, ETH-SLM does not respond to ETH-SLM requests with multicast DA. ETH-SLM for VPLS for point-to-point Ethernet connection is supported using directed unicast destination MAC addresses, although point-to-multipoint topologies are not supported.
- A unicast destination address may be used in provisioned environments for point-to-point connections. However, it requires that the unicast destination address of the downstream MEP must have been configured on the MEP transmitting an alarm indication signal (AIS).
- ETH-SLM is not supported on downstream MEPs on label-switched interfaces (LSIs).
- ETH-SLM is supported on aggregated Ethernet (ae) interfaces
- The number of ETH-SLM sessions for proactive ETH-SLM that can be supported is limited to the total number of iterators that can be supported in the system. This limitation includes the iterator support for other measurement types such as loss, statistical frame loss, and two-way delay. A new iterator type, SLM, is added to support ETH-SLM. The total number of SLA iterators that you can configure in the system is equal to the total number of iterations supported in the system.
- For on-demand SLM, the minimum period between two SLM requests is 100 milliseconds.
- For proactive SLM, the minimum period between two SLM requests is 10 milliseconds for distributed mode and 100 milliseconds for non-distributed mode.
- ETH-SLM frames are always marked as drop-ineligible in compliance with the ITU-T Y.1731 standard.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Transmission of ETH-SLM Messages on page 740
- Format of ETH-SLM Messages on page 738
- Scenarios for Configuration of ETH-SLM on page 737
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Starting an On-Demand ETH-SLM Session on page 747
- Troubleshooting Failures with ETH-SLM on page 752

Starting a Proactive ETH-SLM Session

To start a proactive Ethernet synthetic loss measurement (ETH-SLM) session, you must configure the Ethernet interfaces on maintenance association end points (MEPs) on which packets transmitted with synthetic frame loss need to be analyzed. You must then create an iterator profile to transmit service-level agreement (SLA) measurement packets for ETH-SLM and associate the local and remote MEPs with the profile.

- Configuring MEP Interfaces on page 744
- Configuring an Iterator Profile for ETH-SLM on page 745
- Associating the Iterator Profile with MEPs for ETH-SLM on page 746

Configuring MEP Interfaces

Before you can start an Ethernet synthetic frame loss measurement session across an Ethernet service, you must configure two ACX Series routers to support ETH-SLM.

To configure an Ethernet interface on an ACX Series router to support ETH-SLM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
  ethernet-interface-name {
    vlan-tagging;
    unit logical-unit-number {
      vlan-id vlan-id; # Both interfaces on this VLAN
    }
  }
}
```

   Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md-name { # On both routers
        level number;
        maintenance-association ma-name { # On both routers
          continuity-check {
            interval 100ms;
            hold-interval 1;
          }
          mep mep-id { # Attach to VLAN interface
            auto-discovery;
            direction (up | down);
            interface interface-name;
            priority number;
          }
        }
      }
    }
  }
}
```


Configuring an Iterator Profile for ETH-SLM

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

NOTE: ACX5048 and ACX5096 routers supports iterator cycle time of only 1 second and above.

To create an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   [edit]
   user@host# **edit protocols oam ethernet connectivity-fault-management performance-monitoring**

2. Configure the SLA measurement monitoring iterator:

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# **edit sla-iterator-profiles**

3. Configure an iterator profile—for example, i1:

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
   user@host# **set i1**

4. (Optional) Configure the cycle time, which is the amount of time (in milliseconds) between back-to-back transmission of SLA frames for one connection, with a value from 10 through 3,600,000. The default value is 1000 ms.

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

5. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value), with a value from 1 through 2000. The default value is 2000.

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

6. Configure the measurement type as synthetic loss measurement.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set measurement-type slm**

7. Configure the **disable** statement to stop the iterator (that is, disable the iterator profile).

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set disable**

8. Verify the configuration.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
   user@host# **show i1**
   cycle-time *cycle-time-value*;
   iteration-period *iteration-period-value*;
   measurement-type slm;

Associating the Iterator Profile with MEPs for ETH-SLM

You can associate a remote maintenance association end point (MEP) with more than one iterator profile.

To configure a remote MEP with an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   user@host# **edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id***

2. Configure the remote MEP ID with a value from 1 through 8191.

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

3. Set the iterator profile.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]
   user@host# **set sla-iterator-profile *profile-name***

4. (Optional) Set the size of the data TLV portion of the Y.1731 data frame with a value from 1 through 1400 bytes. The default value is 1.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set data-tlv-size *size***

5. (Optional) Set the iteration count, which indicates the number of iterations for which this connection should partake in the iterator for acquiring SLA measurements, with a value from 1 through 65,535. The default value is 0 (that is, infinite iterations).

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*] user@host# **set** iteration-count ***count-value***

6. (Optional) Set the priority, which is the **vlan-pcp** value that is sent in the Y.1731 data frames, with a value from 0 through 7. The default value is 0.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set** priority ***priority-value***

7. Verify the configuration.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]
   user@host# **show**
   sla-iterator-profile *profile-name* {
     data-tlv-size *size*;
     iteration-count *count-value*;
     priority *priority-value*;
   }

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Transmission of ETH-SLM Messages on page 740
- Format of ETH-SLM Messages on page 738
- Guidelines for Configuring ETH-SLM on page 742
- Scenarios for Configuration of ETH-SLM on page 737
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting an On-Demand ETH-SLM Session on page 747
- Troubleshooting Failures with ETH-SLM on page 752

Starting an On-Demand ETH-SLM Session

To start an on-demand Ethernet synthetic loss measurement (ETH-SLM) session, type the **monitor ethernet synthetic-loss-measurement one-way** command in operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0
      Synthetic Loss measurement statistics:
       SLM packets sent                             : 100
       SLR packets received                         : 100
       Accumulated SLM statistics:
       Local TXFC1 value                            : 100
       Local RXFC1 value                            : 100
```

```
Last Received SLR frame TXFCf(tc)          : 100
Last Received SLR frame TXFCb(tc)          : 100
SLM Frame Loss:
Frame Loss (far-end)                       : 0 (0.00 %)
Frame Loss (near-end)                      : 0 (0.00 %)
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must press **Ctrl + C** to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Transmission of ETH-SLM Messages on page 740
- Format of ETH-SLM Messages on page 738
- Guidelines for Configuring ETH-SLM on page 742
- Scenarios for Configuration of ETH-SLM on page 737
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Troubleshooting Failures with ETH-SLM on page 752

Managing ETH-SLM Statistics and ETH-SLM Frame Counts

Displaying ETH-SLM Statistics Only

Purpose Display on-demand ETH-SLM statistics.

By default, the **show oam ethernet connectivity-fault-management synthetic-loss-statistics** command displays on-demand ETH-SLM statistics for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

Action

- To display the on-demand ETH-SLM statistics collected for MEPs belonging to maintenance association **ma1** within maintenance domain **md1**:

user@host> **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md1 maintenance-association ma1**

- To display the on-demand ETH-SLM statistics collected for ETH-SLM sessions for the local MEP **201** belonging to maintenance association **ma2** within maintenance domain **md2**:

  user@host> **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the on-demand ETH-SLM statistics collected for ETH-SLM sessions from local MEPs belonging to maintenance association **ma3** within maintenance domain **md3** to the remote MEP **302**:

  user@host> **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md3 maintenance-association ma3 remote-mep 302**

Meaning The output displays on-demand ETH-SLM statistics for MEPs in the specified maintenance association within the specified maintenance domain. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management synthetic-loss-statistics**.

See Also

- **show oam ethernet connectivity-fault-management synthetic-loss-statistics**

Displaying ETH-SLM Statistics and Frame Counts

Purpose Display on-demand ETH-SLM statistics and ETH-SLM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays on-demand ETH-SLM statistics and frame counts for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

Action

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for MEPs in maintenance association **ma1** within maintenance domain **md1**:

  user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1**

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for the local MEP **201** in maintenance association **ma2** within maintenance domain **md2**:

  user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for the local MEP in maintenance association **ma3** within maintenance domain **md3** that participates in an ETH-SLM session with the remote MEP **302**:

`user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Meaning The output displays on-demand ETH-SLM statistics and ETH-SLM frame counts for MEPs in the specified maintenance association within the specified maintenance domain. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management mep-statistics**.

See Also

- **show oam ethernet connectivity-fault-management mep-statistics**

Displaying ETH-SLM Frame Counts for MEPs by Enclosing CFM Entity

Purpose Display on-demand ETH-SLM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display the ETH-SLM frame counts and not the MEP database details.

Action

- To display CFM database information (including ETH-SLM frame counts) for all MEPs in MA **ma1** within maintenance domain **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma1 maintenance-association ma1**

- To display CFM database information (including ETH-SLM frame counts) only for the local MEP **201** in MA **ma1** within maintenance domain **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display CFM database information (including ETH-SLM frame counts) only for the remote MEP **302** in MA **ma3** within maintenance domain **md3**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Meaning The output displays ETH-SLM frame counts for MEPs within a particular maintenance domain, or for a specific local or remote MEP. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management mep-database**.

See Also
- **show oam ethernet connectivity-fault-management mep-database**

Displaying ETH-SLM Frame Counts for MEPs by Interface or Domain Level

Purpose Display on-demand ETH-SLM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-SLM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

NOTE: **At the router attached to the initiator MEP, you can only display the ETH-SLM frame counts and not the MEP database details.**

Action
- To display CFM database information (including ETH-SLM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces detail**

- To display CFM database information (including ETH-SLM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail**

- To display CFM database information (including ETH-SLM frame counts) only for MEPs enclosed within CFM maintenance domains at level **6**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces level 6 detail**

Meaning The output displays ETH-SLM frame counts for MEPs for the specified interface. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management interfaces**.

See Also
- **show oam ethernet connectivity-fault-management interfaces**

Clearing ETH-SLM Statistics and Frame Counts

Purpose Clear the on-demand ETH-SLM statistics and ETH-SLM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the on-demand ETH-SLM statistics and ETH-SLM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management synthetic-loss-measurement**

- To clear the on-demand ETH-SLM statistics and ETH-SLM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

  `user@host>` **clear oam ethernet connectivity-fault-management synthetic-loss-measurement ge-0/5/9.0**

Clearing Iterator Statistics

Purpose

Clear the existing iterator statistics and proactive ETH-SLM counters.

Multiple iterators can be associated with remote MEP. However, by default, only one result pertaining to one iterator profile can be cleared.

Action

- To clear the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep *1***

- To clear the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1**

Related Documentation

- clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 1471
- show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 2088
- show oam ethernet connectivity-fault-management interfaces on page 2037 (**detail** | **extensive**)
- show oam ethernet connectivity-fault-management mep-statistics on page 2066
- show oam ethernet connectivity-fault-management mep-database on page 2055

Troubleshooting Failures with ETH-SLM

Problem

Description: The Ethernet synthetic loss measurement (ETH-SLM) application is not working properly for calculation of frame loss using synthetic frames instead of data traffic

Solution Perform the following steps to analyze and debug any problems with the ETH-SLM functionality.

1. Ensure that ETH-SLM is configured (either proactive or on-demand) to initiate SLM frames. Verify the configuration settings.

2. Examine any failures that might have occurred in the CFM session for which the ETH-SLM feature is enabled. The CFM session must be in the up state for the ETH-SLM functionality to work correctly. Use the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*** **local-mep** ***mep-id*** **remote-mep** ***remote-mep-id*** command to verify whether the CFM session is in the up state.

3. If the MEP sessions are active, use the appropriate show command to verify the ETH-SLM statistics and to analyze if ETH-SLM frames are transmitted or received.

4. If the transmission of ETH-SLM frames does not happen correctly after you attempt all of the preceding troubleshooting steps, enable the tracing operations for Ethernet CFM by including the **traceoptions** statement at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
traceoptions {
  file <filename> <files number  <match regular-expression microsecond-stamp>>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 735
- Transmission of ETH-SLM Messages on page 740
- Format of ETH-SLM Messages on page 738
- Guidelines for Configuring ETH-SLM on page 742
- Scenarios for Configuration of ETH-SLM on page 737
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 748
- Starting a Proactive ETH-SLM Session on page 743
- Starting an On-Demand ETH-SLM Session on page 747

Configuring an Iterator Profile

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

To create an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   [edit]
   user@host# **edit protocols oam ethernet connectivity-fault-management performance-monitoring**

2. Configure the SLA measurement monitoring iterator:

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# **edit** sla-iterator-profiles

3. Configure an iterator profile—for example, i1:

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
   user@host# **set i1**

4. (Optional) Configure the cycle time, which is the amount of time (in milliseconds) between back-to-back transmission of SLA frames for one connection, with values from 10 through 3,600,000. The default value is 1000 ms.

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

5. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value), with values from 1 through 2000. The default value is 2000.

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

6. Configure the measurement type as loss measurement, statistical frame-loss measurement, or two-way delay measurement.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set** measurement-type **(loss | statistical-frame-loss | two-way-delay)**

7. (Optional) Configure the calculation weight for delay with values from 1 through 65,535. The default value is 1 (applicable only for two-way delay measurement).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
   sla-iterator-profiles i1]
user@host# set calculation-weight delay delay-value
```

8. (Optional) Configure the calculation weight for delay variation with values from 1 through 65,535. The default value is 1 (applicable only for two-way delay measurement).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
   sla-iterator-profiles i1]
user@host# set calculation-weight delay-variation delay-variation-value
```

9. (Optional) Configure the threshold value for average frame delay, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay is exceeded, an SNMP trap is generated for ETH-DM. The range is from 1 through 4294967295 microseconds.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
   sla-iterator-profiles i1]
user@host# set avg-fd-twoway-threshold avg-fd-twoway-threshold-value
```

10. (Optional) Configure the threshold value for average frame delay variation, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay variationis exceeded, an SNMP trap is generated for ETH-DM. The range is from 1 through 4294967295 microseconds.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
   sla-iterator-profiles i1]
user@host# set avg-ifdv-twoway-threshold avg-ifdv-twoway-threshold-value
```

11. (Optional) Configure the threshold value for average frame loss ratio, in milli-percent, in the upward or forward direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average forward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. The range is from 1 through 100000 milli-percent.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
   sla-iterator-profiles i1]
user@host# set avg-flr-forward-threshold avg-flr-forward-threshold-value
```

12. (Optional) Configure the threshold value for average frame loss ratio, in milli-percent, in the backward or downstream direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average backward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. The range is from 1 through 100000 milli-percent.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
   sla-iterator-profiles i1]
user@host# set avg-flr-backward-threshold avg-flr-backward-threshold-value
```

13. Configure the **disable** statement to stop the iterator (that is, disable the iterator profile).

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

14. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
  sla-iterator-profiles]
user@host# show i1
  cycle-time cycle-time-value;
  iteration-period iteration-period-value;
  measurement-type (loss | two-way-delay);
  avg-fd-twoway-threshold avg-fd-twoway-threshold-value;
  avg-ifdv-twoway-threshold avg-ifdv-twoway-threshold-value;
  avg-flr-forward-threshold avg-flr-forward-threshold-value;
  avg-flr-backward-threshold avg-flr-backward-threshold-value;
  calculation-weight {
    delay delay-weight;
    delay-variation delay-variation-weight;
  }
  calculation-weight {
    delay delay-weight;
    delay-variation delay-variation-weight;
  }
```

Related Documentation

- Proactive Mode for SLA Measurement on page 733
- Configuring a Remote MEP with an Iterator Profile on page 765
- Verifying the Configuration of an Iterator Profile on page 756
- Managing Iterator Statistics on page 759

Verifying the Configuration of an Iterator Profile

The following topics illustrate the configuration of an iterator profile for a two-way delay measurement, for loss measurement, and for a remote maintenance association end point (MEP). The topics also illustrate disabling an iterator profile with the **disable** statement for two-way measurement and deactivating an iterator profile with the **deactivate** command for a remote MEP.

- Displaying the Configuration of an Iterator Profile for Two-way Delay Measurement on page 757
- Displaying the Configuration of an Iterator Profile for Loss Measurement on page 757
- Displaying the Configuration of a Remote MEP with an Iterator Profile on page 758
- Disabling an Iterator Profile on page 759

Displaying the Configuration of an Iterator Profile for Two-way Delay Measurement

Purpose Display the configuration of an iterator profile for two-way delay measurement as configured in the “Configuring an Iterator Profile” on page 754 topic with the following values:

- **profile-name—i1**
- **cycle-time—1000** milliseconds
- **iteration-period—2000** cycles per second
- **delay—1**
- **delay-variation—1**:

Action To display information about the iterator profile, run the **show** command at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
  sla-iterator-profiles]
user@host# show
i1 {
  cycle-time 1000;
  iteration-period 2000;
  measurement-type two-way-delay;
  calculation-weight {
    delay 1;
    delay-variation 1;
  }
}
```

Meaning The configuration for an iterator profile for two-way measurement is displayed as expected with set values.

Displaying the Configuration of an Iterator Profile for Loss Measurement

Purpose Display the configuration of an iterator profile for loss measurement as configured in the “Configuring an Iterator Profile” on page 754 topic with the following values:

- **profile-name—12**
- **cycle-time—1000** milliseconds
- **iteration-period—2000** cycles per second

Action To display information about the iterator profile, run the **show** command at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]** hierarchy level:


```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
  sla-iterator-profiles]
user@host# show
12 {
  cycle-time 1000;
  iteration-period 2000;
  measurement-type loss;
}
```

Meaning The configuration for an iterator profile for loss measurement is displayed as expected with set values.

Displaying the Configuration of a Remote MEP with an Iterator Profile

Purpose Display the configuration of a remoteMEP as configured in the “Configuring a Remote MEP with an Iterator Profile” on page 765 topic with the following values:

- **profile-name—i3**
- **maintenance-domain—default-1**
- **maintenance-association—1**
- **short-name-format—2octet**
- **mep—1**
- **remote-mep—1**
- **data-tlv-size—1**
- **iteration-count—1**
- **priority—1**

Action To display information about the remote MEP, run the **show** command at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain default-1 maintenance association ma1 mep 1 remote-mep 1]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  default-1 maintenance association 1 short-name-format 2octet mep 1 remote-mep 1]
user@host# show
sla-iterator-profile i3 {
  data-tlv-size 1;
  iteration-count 1;
  priority 1;
}
```

Meaning The configuration for a remote MEP for two-way measurement is displayed as expected with set values.

Disabling an Iterator Profile

Purpose To disable an iterator profile for two-way delay measurement and for a remote MEP.

Action

- To disable an iterator profile (for example, i1) with the **disable** configuration command for two-way measurement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]** hierarchy level:

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

- To disable an iterator profile for a remote MEP (for example, i2) with the **deactivate** configuration command at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain default-1 maintenance association ma1 mep 1 remote-mep 1]** hierarchy level:

  ```
  [edit protocols oam ethernet connectivity-fault-management maintenance-domain
    default-1 maintenance association ma1 mep 1 remote-mep 1]
  user@host# deactivate sla-iterator-profile i2
  ```

Related Documentation

- Proactive Mode for SLA Measurement on page 733
- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Managing Iterator Statistics on page 759

Managing Iterator Statistics

- Displaying Iterator Statistics on page 759
- Clearing Iterator Statistics on page 764

Displaying Iterator Statistics

Purpose Retrieve and display iterator statistics.

Muliple iterators can be associated with a remote MEP. However, by default, only one result pertaining to one iterator profile is displayed.

Action

- To display the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1** (here, the iterator profile **i1** is configured for two-way delay measurement):

  ```
  user@host> show oam ethernet connectivity-fault-management sla-iterator-statistics
  sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1
  remote-mep 1
  ```

```
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent                                       : 23898520
DMM skipped for threshold hit                  : 11000
DMM skipped for threshold hit window           : 0
DMR received                                   : 23851165
DMR out of sequence                            : 1142
DMR received with invalid time stamps          : 36540
Average two-way delay                          : 129 usec
Average two-way delay variation                : 15 usec
Average one-way forward delay variation        : 22 usec
Average one-way backward delay variation       : 22 usec
Weighted average two-way delay                 : 134 usec
Weighted average two-way delay variation       : 8 usec
Weighted average one-way forward delay variation  : 6 usec
Weighted average one-way backward delay variation : 2 usec
```

Output fields are listed in the approximate order in which they appear.

Table 71: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields

Output Field Name	Output Field Description
Maintenance domain	Maintenance domain name.
Level	Maintenance domain level configured.
Maintenance association	Maintenance association name.
Local MEP id	Numeric identifier of the local MEP.
Remote MEP id	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.

Table 71: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields (continued)

Output Field Name	Output Field Description
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
Delay weight	Calculation weight of delay.
Delay variation weight	Calculation weight of delay variation.
DMM sent	Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.
DMM skipped for threshold hit	Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.
DMM skipped for threshold hit window	Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
DMR received	Number of delay measurement reply (DMR) frames received.
DMR out of sequence	Total number of DMR out of sequence packets received.
DMR received with invalid time stamps	Total number of DMR frames received with invalid timestamps.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Average one-way forward delay variation	Average one-way forward delay variation for the statistics displayed in microseconds.
Average one-way backward delay variation	Average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average two-way delay	Weighted average two-way delay for the statistics displayed in microseconds.
Weighted average two-way delay variation	Weighted average two-way delay variation for the statistics displayed in microseconds.
Weighted average one-way forward delay variation	Weighted average one-way forward delay variation for the statistics displayed in microseconds.
Weighted average one-way backward delay variation	Weighted average one-way backward delay variation for the statistics displayed in microseconds.

- To display the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1** (here, the iterator profile **i1** is configured for loss measurement):

```
user@host> show oam ethernet connectivity-fault-management sla-iterator-statistics
sla-iterator i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1
remote-mep 1
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i2, Iterator Id: 2
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:25 ago)
Reset reason: Adjacency flap

Iterator loss measurement statistics:
LMM sent                              : 238970
LMM skipped for threshold hit         : 60
LMM skipped for threshold hit window  : 0
LMR received                          : 238766
LMR out of sequence                   : 43

Accumulated transmit statistics:
Near-end (CIR)                        : 0
Far-end (CIR)                         : 0
Near-end (EIR)                        : 0
Far-end (EIR)                         : 0

Accumulated loss statistics:
Near-end (CIR)                        : 0 (0.00%)
Far-end (CIR)                         : 0 (0.00%)
Near-end (EIR)                        : 0 (0.00%)
Far-end (EIR)                         : 0 (0.00%)

Last loss measurement statistics:
Near-end (CIR)                        : 0
Far-end (CIR)                         : 0
Near-end (EIR)                        : 0
Far-end (EIR)                         : 0
```

Output fields are listed in the approximate order in which they appear.

Table 72: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields

Output Field Name	Output Field Description
Maintenance domain	Maintenance domain name.
Level	Maintenance domain level configured.
Maintenance association	Maintenance association name.
Local MEP id	Numeric identifier of the local MEP.
RemoteMEP identifier	Numeric identifier of the remote MEP.

Table 72: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields (continued)

Output Field Name	Output Field Description
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
LMM sent	Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.
LMM skipped for threshold hit	Number of LMM frames sent to the peer MEP in this session skipped during threshold hit.
LMM skipped for threshold hit window	Number of LMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
LMR received	Number of LMRs frames received.
LMR out of sequence	Total number of LMR out of sequence packets received.
Near-end (CIR)	Frame loss associated with ingress data frames for the statistics displayed.
Far-end (CIR)	Frame loss associated with egress data frames for the statistics displayed.
Near-end (EIR)	Frame loss associated with ingress data frames for the statistics displayed.
Far-end (EIR)	Frame loss associated with egress data frames for the statistics displayed.

See Also

- Proactive Mode for SLA Measurement on page 733
- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Verifying the Configuration of an Iterator Profile on page 756

- *Ethernet Interfaces Feature Guide for Routing Devices*

Clearing Iterator Statistics

Purpose

Clear iterator statistics.

Multiple iterators can be associated with remote MEP. However, by default, only one result pertaining to one iterator profile can be cleared.

Action

- To clear the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep** ***1***

- To clear the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1**

See Also

- Proactive Mode for SLA Measurement on page 733
- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Verifying the Configuration of an Iterator Profile on page 756
- Proactive Mode for SLA Measurement on page 733

- *Ethernet Interfaces Feature Guide for Routing Devices*

Related Documentation

- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Verifying the Configuration of an Iterator Profile on page 756
- Proactive Mode for SLA Measurement on page 733

Configuring a Remote MEP with an Iterator Profile

You can associate a remote maintenance association end point (MEP) with more than one iterator profile.

To configure a remote MEP with an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   user@host# **edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*** mep ***mep-id***

2. Configure the remote MEP with values from 1 through 8191.

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

3. Set the iterator profile.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]
   user@host# **set** sla-iterator-profile ***profile-name***

4. (Optional) Set the size of the data TLV portion of the Y.1731 data frame with values from 1 through 1400 bytes. The default value is 1.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set** data-tlv-size ***size***

5. (Optional) Set the iteration count, which indicates the number of iterations for which this connection should partake in the iterator for acquiring SLA measurements, with values from 1 through 65,535. The default value is 0 (that is, infinite iterations).

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set** iteration-count ***count-value***

6. (Optional) Set the priority, which is the **vlan-pcp** value that is sent in the Y.1731 data frames, with values from 0 through 7. The default value is 0.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set** priority ***priority-value***

7. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name mep mep-id remote-mep
  remote-mep-id]
user@host# show
sla-iterator-profile profile-name {
  data-tlv-size size;
  iteration-count count-value;
  priority priority-value;
}
```

Related Documentation

- Proactive Mode for SLA Measurement on page 733
- Configuring an Iterator Profile on page 754
- Verifying the Configuration of an Iterator Profile on page 756
- Managing Iterator Statistics on page 759

Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS

You can dampen the performance monitoring threshold-crossing traps and notifications that are generated every time a threshold-crossing event occurs to prevent congestion of the network management system (NMS).

Damping limits the number of jnxSoamPmThresholdCrossingAlarm traps sent to the NMS by summarizing the flap occurrences over a period of time, known as the flap trap timer, and sends a single jnxSoamPmThresholdFlapAlarm notification to the NMS. You can configure the duration of the flap trap timer to any value from 1 through 360 seconds.

The jnxSoamPmThresholdFlapAlarm notification is generated and sent when the following conditions are met:

- At least one flap has occurred when the flap timer has expired.
- You changed the value of the flap trap timer, which caused the timer to stop.

You can enable damping at the global level for the iterator or you can enable damping at the individual threshold type of the iterator. For instance, to enable damping at the global level, for the iterator, use the following command: **set protocols oam ethernet cfm performance-monitoring sla-iterator-profiles** ***profile-name*** **flap-trap-monitor**. To enable damping at a specific threshold type, for the **avg-fd-twoway-threshold**, use the following command: **set protocols oam ethernet cfm performance-monitoring sla-iterator-profiles** ***profile-name*** **avg-fdv-twoway-threshold flap-trap-monitor**.

You can also disable damping.

Related Documentation

- flap-trap-monitor on page 1013

Configuring Statistical Frame Loss Measurement for VPLS Connections

Using proactive statistical frame loss measurement, you can monitor VPLS connections on MX Series routers . Statistical frame loss measurement allows you to monitor the quality of Ethernet connections for service level agreements (SLAs). Point-to-point and multipoint-to-multipoint connections configured on MX Series routers can be monitored by registering the connection on an iterator and initiating periodic SLA measurement of frame transmissions on the connections.

Iterators periodically transmit SLA measurement packets using ITU-Y.1731 compliant frames. The iterator sends periodic measurement packets for each of the connections registered to it. These measurement cycles are transmitted in such a way as to not overlap, reducing the processing demands placed on the CPU. The measurement packets are exchanged between the source user network interface (UNI) port and the destination UNI port, providing a sequence of timed performance measurements for each UNI pair. The Frame Loss Ratio (FLR) and connection availability can be computed from these measurements using statistics.

The following steps outline how to configure statistical frame loss measurement for VPLS connections:

1. To configure proactive ETH-DM measurement for a VPLS connection, see "Guidelines for Configuring Routers to Support an ETH-DM Session" on page 768.

2. To enable statistical loss measurement for a VPLS connection, configure an iterator for the VPLS connection using the **sla-iterator-profiles** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level. For detailed instructions, see "Configuring an Iterator Profile" on page 754.

3. As part of the iterator configuration, include the **statistical-frame-loss** option for the **measurement-type** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]** hierarchy level.

4. Once you have enabled the iterator, you can display the statistical frame loss for a VPLS connection by issuing the **show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator** ***identifier*** **maintenance-domain** ***name*** **maintenance-association** ***name*** **local-mep** ***identifier*** **remote-mep** ***identifier*** command.

Related Documentation

- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768
- Configuring an Iterator Profile on page 754
- Verifying the Configuration of an Iterator Profile on page 756

Guidelines for Configuring Routers to Support an ETH-DM Session

Keep the following guidelines in mind when configuring routers to support an Ethernet frame delay measurement (ETH-DM) session:

- Configuration Requirements for ETH-DM on page 768
- Configuration Options for ETH-DM on page 768

Configuration Requirements for ETH-DM

You can obtain ETH-DM information for a link that meets the following requirements:

- The measurements can be performed between peer maintenance association endpoints (MEPs) on two routers.
- The two MEPs must be configured on two Ethernet physical interfaces or on two Ethernet logical interfaces. For more information, see "Configuring a MEP to Generate and Respond to CFM Protocol Messages" on page 613.
- The two MEPs must be configured—on their respective routers—under the same maintenance association (MA) identifier. For more information, see "Creating a Maintenance Association" on page 610.
- On both routers, the MA must be associated with the same maintenance domain (MD) name. For more information, see "Creating a Maintenance Domain" on page 603.
- On both routers, periodic packet management (PPM) must be running on the Routing Engine and Packet Forwarding Engine, which is the default configuration. You can disable PPM on the Packet Forwarding Engine only. However, the Ethernet frame delay measurement feature requires that distributed PPM remain enabled on the Packet Forwarding Engine of both routers. For more information about **ppm**, see the *Junos OS Routing Protocols Library*.
- If the PPM process (**ppm**) is disabled on the Packet Forwarding Engine, you must re-enable it. Re-enabling distributed **ppm** entails restarting the **ethernet-connectivity-fault-management** process, which causes all connectivity fault management (CFM) sessions to re-establish. For more information about CFM sessions, see "Configuring Ethernet Local Management Interface" on page 627.

NOTE: **The Ethernet frame delay measurement feature is supported only for MEPs configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.**

Configuration Options for ETH-DM

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Configuring Routers to Support an ETH-DM Session on page 775

Guidelines for Starting an ETH-DM Session

Keep the following guidelines in mind when preparing to start an Ethernet frame delay measurement (ETH-DM) session:

- ETH-DM Session Prerequisites on page 769
- ETH-DM Session Parameters on page 769
- Restrictions for an ETH-DM Session on page 770

ETH-DM Session Prerequisites

Before you can start an ETH-DM session, you must configure two MX Series routers to support ETH-DM by defining the two CFM-enabled physical or logical Ethernet interfaces on each router. This entails creating and configuring CFM maintenance domains, maintenance associations, and maintenance association end points on each router. For more information about enabling CFM on an Ethernet interface, see "Creating a Maintenance Domain" on page 603.

NOTE: **The Ethernet frame delay measurement feature is supported only for maintenance association end points configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.**

For specific information about configuring routers to support ETH-DM, see "Guidelines for Configuring Routers to Support an ETH-DM Session" on page 768 and "Configuring Routers to Support an ETH-DM Session" on page 775.

ETH-DM Session Parameters

You can initiate a one-way or two-way ETH-DM session by entering the **monitor ethernet delay-measurement** operational command at a router that contains one end of the service for which you want to measure frame delay. The command options specify the ETH-DM session in terms of the CFM elements:

- The type of ETH-DM measurement (one-way or two-way) to be performed.
- The Ethernet service for which the ETH-DM measurement is to be performed:
  - CFM maintenance domain—Name of the existing maintenance domain (MD) for which you want to measure Ethernet frame delays. For more information, see "Creating a Maintenance Domain" on page 603.

  - CFM maintenance association—Name of an existing maintenance association (MA) within the maintenance domain. For more information, see "Creating a Maintenance Association" on page 610.
  - Remote CFM maintenance association end point—The unicast MAC address or the numeric identifier of the remote maintenance association end point (MEP)—the physical or logical interface on the remote router that resides in the specified MD and is named in the specified MA—with which to perform the ETH-DM session. For more information, see "Configuring a MEP to Generate and Respond to CFM Protocol Messages" on page 613.

- Optional specifications:
  - Count—You can specify the number of ETH-DM requests to send for this frame delay measurement session. The range is from 1 through 65,535 frames. The default value is 10 frames.

    NOTE: Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).

  - Frame interval—You can specify the number of seconds to elapse between ETH-DM frame transmittals. The default value is 1 second.

For more detailed information about the parameters you can specify to start an ETH-DM session, see the **monitor ethernet delay-measurement** operational command description in the CLI Explorer.

Restrictions for an ETH-DM Session

The following restrictions apply to an ETH-DM session:

- You cannot run multiple simultaneous ETH-DM sessions with the same remote MEP or MAC address.
- For a given ETH-DM session, you can collect frame delay information for a maximum of 65,535 frames.
- For a given CFM session (pair of peer MEPs), the ETH-DM database stores a maximum of 100 statistics, with the older statistics being "aged out" as newer statistics are collected for that pair of MEPs.
  - For one-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the receiver MEP is defined.
  - For two-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the initiator MEP is defined.

  Depending on the number of frames exchanged in the individual ETH-DM sessions, the ETH-DM database can contain statistics collected through multiple ETH-DM sessions.

- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *Junos OS High Availability Library for Routing Devices*.
- Accuracy of frame delay data is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Starting an ETH-DM Session on page 780
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 771

- **monitor ethernet delay-measurement** operational command

Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts

ETH-DM Statistics

Ethernet frame delay statistics are the frame delay and frame delay variation values determined by the exchange of frames containing ETH-DM protocol data units (PDUs).

- For a one-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, including the exchange of one-way delay PDU frames, see "Ethernet Frame Delay Measurements Overview" on page 723.
- For a two-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, including the exchange of two-way delay PDU frames, see "Ethernet Frame Delay Measurements Overview" on page 723.

A CFM database stores CFM-related statistics and—for Ethernet interfaces that support ETH-DM—the 100 most recently collected ETH-DM statistics for that pair of MEPs. You can view ETH-DM statistics by using the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command to display the CFM statistics for the MEP that collects the ETH-DM statistics you want to view.

Table 73 on page 772 describes the ETH-DM statistics calculated in an ETH-DM session.

Table 73: ETH-DM Statistics

Field Name	Field Description
One-way delay (μsec)†	For a one-way ETH-DM session, the frame delay, in microseconds, collected at the receiver MEP. To display frame delay statistics for a given one-way ETH-DM session, use the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command at the receiver MEP for that session.
Two-way delay (μsec)	For a two-way ETH-DM session, the frame delay, in microseconds, collected at the initiator MEP. When you start a two-way frame delay measurement, the CLI output displays each DMR frame receipt timestamp and corresponding DMM frame delay and delay variation collected as the session progresses. To display frame delay statistics for a given two-way ETH-DM session, use the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command at the initiator MEP for that session.
Average delay†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Average delay** field displays the average one-way and two- frame delays among all ETH-DM statistics collected at the CFM session level. For example, suppose you start two one-way ETH-DM sessions for 50 counts each, one after the other. If, after both measurement sessions complete, you use a **show** command to display 100 ETH-DM statistics for that CFM session, the **Average delay** field displays the average frame delay among all 100 statistics.
Average delay variation†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay variation among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Average delay variation** field displays the average one-way and two- frame delay variations among all ETH-DM statistics collected at the CFM session level.
Best-case delay†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the lowest two-way frame delay value among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Best case delay** field displays the lowest one-way and two-way frame delays among all ETH-DM statistics collected at the CFM session level.
Worst-case delay†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the highest two-way frame delay value among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Worst case delay** field displays the highest one-way and two-way frame delays among all statistics collected at the CFM session level.

Table 73: ETH-DM Statistics (continued)

Field Name	Field Description
[†]When you start a one-way frame delay measurement, the CLI output displays **NA** (“not available”) for this field. One-way ETH-DM statistics are collected at the remote (receiver) MEP. Statistics for a given one-way ETH-DM session are available only by displaying CFM statistics for the receiver MEP.	

ETH-DM Statistics Retrieval

At the receiver MEP for a one-way session, or at the initiator MEP for a two-way session, you can display all ETH-DM statistics collected at a CFM session level by using the following operational commands:

- **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*> <count *count*>**
- **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*> <count *count*>**

ETH-DM Frame Counts

The number of ETH-DM PDU frames exchanged in a ETH-DM session are stored in the CFM database on each router.

Table 74 on page 773 describes the ETH-DM frame counts collected in an ETH-DM session.

Table 74: ETH-DM Frame Counts

Field Name	Field Description
1DMs sent	Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. Stored in the CFM database of the MEP initiating a one-way frame delay measurement.
Valid 1DMs received	Number of valid 1DM frames received. Stored in the CFM database of the MEP receiving a one-way frame delay measurement.
Invalid 1DMs received	Number of invalid 1DM frames received. Stored in the CFM database of the MEP receiving a one-way frame delay measurement.
DMMs sent	Number of delay measurement message (DMM) PDU frames sent to the peer MEP in this session. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.
DMRs sent	Number of delay measurement reply (DMR) frames sent (in response to a received DMM). Stored in the CFM database of the MEP responding to a two-way frame delay measurement.

Table 74: ETH-DM Frame Counts (continued)

Field Name	Field Description
Valid DMRs received	Number of valid DMR frames received. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.
Invalid DMRs received	Number of invalid DMR frames received. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.

ETH-DM Frame Count Retrieval

Each router counts the number of ETH-DM frames sent or received and stores the counts in a CFM database.

Frame Counts Stored in CFM Databases

You can display ETH-DM frame counts for MEPs assigned to specified Ethernet interfaces or for specified MEPs in CFM sessions by using the following operational commands:

- **show oam ethernet connectivity-fault-management interfaces** (**detail** | **extensive**)
- **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*** <**local-mep** ***mep-id***> <**remote-mep** ***mep-id***>

One-Way ETH-DM Frame Counts

For a one-way ETH-DM session, delay statistics are collected at the receiver MEP only, but frame counts are collected at both MEPs. As indicated in Table 74 on page 773, one-way ETH-DM frame counts are tallied from the perspective of each router in the session:

- At the initiator MEP, the router counts the number of 1DM frames sent.
- At the receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

You can also view one-way ETH-DM frame counts—for a receiver MEP—by using the **show oam ethernet connectivity-fault-management mep-statistics** command to display one-way statistics and frame counts together.

Two-Way ETH-DM Frame Counts

For a two-way ETH-DM session, delay statistics are collected at the initiator MEP only, but frame counts are collected at both MEPs. As indicated in Table 74 on page 773, two-way ETH-DM frame counts are tallied from the perspective of each router in the session:

- At the initiator MEP, the router counts the number of DMM frames sent, valid DMR frames received, and invalid DMR frames received.
- At the responder MEP, the router counts the number of DMR frames sent.

You can also view two-way ETH-DM frame counts—for an initiator MEP—by using the **show oam ethernet connectivity-fault-management mep-statistics** command to display two-way statistics and frame counts together.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 784
- **clear oam ethernet connectivity-fault-management statistics on page 1469** command
- **show oam ethernet connectivity-fault-management mep-statistics on page 2066** command
- **show oam ethernet connectivity-fault-management delay-statistics on page 2029** command
- **show oam ethernet connectivity-fault-management interfaces on page 2037** (**detail** | **extensive**) command
- **show oam ethernet connectivity-fault-management mep-database on page 2055** command

Configuring Routers to Support an ETH-DM Session

Configuring MEP Interfaces

Before you can start an Ethernet frame delay measurement session across an Ethernet service, you must configure two MX Series routers to support ETH-DM.

To configure an Ethernet interface on a MX Series router to support ETH-DM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
   ethernet-interface-name {
      vlan-tagging;
      unit logical-unit-number {
         vlan-id vlan-id; # Both interfaces on this VLAN
      }
   }
}
```

   Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```
[edit protocols]
oam {
   ethernet {
```

```
        connectivity-fault-management {
          maintenance-domain md-name { # On both routers
            level number;
            maintenance-association ma-name { # On both routers
              continuity-check {
                interval 100ms;
                hold-interval 1;
              }
              mep mep-id { # Attach to VLAN interface
                auto-discovery;
                direction (up | down);
                interface interface-name;
                priority number;
              }
            }
          }
        }
      }
    }
```

See Also

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Ensuring That Distributed ppm Is Not Disabled

By default, the router's period packet management process (**ppm**) runs sessions distributed to the Packet Forwarding Engine in addition to the Routing Engine. This process is responsible for periodic transmission of packets on behalf of its various client processes, such as Bidirectional Forwarding Detection (BFD), and it also receives packets on behalf of client processes.

In addition, **ppm** handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With **ppm** processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run such processes as BFD on the Packet Forwarding Engine.

Distributed ppm Required for ETH-DM

Ethernet frame delay measurement requires that **ppm** remains distributed to the Packet Forwarding Engine. If **ppm** is not distributed to the Packet Forwarding Engines of both routers, ETH-DM PDU frame timestamps and ETH-DM statistics are not valid.

Before you start ETH-DM, you must verify that the following configuration statement is *NOT* present:

```
[edit]
routing-options {
  ppm {
    no-delegate-processing;
  }
}
```

If distributed **ppm** processing is disabled (as shown in the stanza above) on either router, you must re-enable it in order to use the ETH-DM feature.

Procedure to Ensure that Distributed ppm is Not Disabled

To ensure that distributed **ppm** is not disabled on a router:

1. Display the packet processing management (PPM) configuration to determine whether distributed **ppm** is disabled.

   - In the following example, distributed **ppm** is enabled on the router. In this case, you do not need to modify the router configuration:

     ```
     [edit]
     user@host# show routing-options
     ppm;
     ```

   - In the following example, distributed **ppm** is disabled on the router. In this case, you must proceed to Step 2 to modify the router configuration:

     ```
     [edit]
     user@host show routing-options
     ppm {
       no-delegate-processing;
     }
     ```

2. Modify the router configuration to re-enable distributed **ppm** and restart the Ethernet OAM Connectivity Fault Management process *ONLY IF* distributed **ppm** is disabled (as determined in the previous step).

   a. Before continuing, make any necessary preparations for the possible loss of connectivity on the router.

      Restarting the **ethernet-connectivity-fault-management** process has the following effect on your network:

      - All connectivity fault management (CFM) sessions re-establish.
      - All ETH-DM requests on the router terminate.
      - All ETH-DM statistics and frame counts reset to 0.

   b. Modify the router configuration to re-enable distributed **ppm**. For example:

      ```
      [edit]
      user@host# delete routing-options ppm no-delegate-processing
      ```

   c. Commit the updated router configuration. For example:

      ```
      [edit]
      user@host# commit and-quit
      commit complete
      exiting configuration mode
      ```

   d. To restart the Ethernet OAM Connectivity-Fault-Management process, enter the **restart ethernet-connectivity-fault-management <gracefully | immediately | soft>** operational mode command. For example:

      ```
      user@host> restart ethernet-connectivity-fault-management
      Connectivity fault management process started, pid 9893
      ```

Connectivity fault management (CFM) sessions operate in centralized mode over AE interfaces by default. Y.1731 performance monitoring (PM) is supported on centralized CFM sessions over AE interfaces. Also, distribution of CFM session over AE interfaces to line cards is supported from Junos OS Release 13.3. To enable the distribution of CFM sessions and to operate in centralized mode, include the **ppm delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level. The mechanism that enables distribution of CFM sessions over AE interfaces provides the underlying infrastructure to support PM over AE interfaces. In addition, periodic packet management (PPM) handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With PPM processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run performance monitoring processes on the Packet Forwarding Engine.

See Also

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Enabling the Hardware-Assisted Timestamping Option

By default, Ethernet frame delay measurement uses software for timestamping transmitted and received ETH-DM frames. For Ethernet interfaces, you can optionally use hardware timing to assist in the timestamping of received ETH-DM frames to increase the accuracy of delay measurements.

Enabling hardware-assisted timestamping of received frames can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction.

To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      performance-monitoring {
        hardware-assisted-timestamping;
      }
    }
  }
}
```

See Also

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Configuring the Server-Side Processing Option

You can delegate the server-side processing (for both two-way delay measurement and loss measurement) to the Packet Forwarding Engine to prevent overloading on the Routing Engine. By default, the server-side processing is done by the Routing Engine.

To configure the server-side processing option:

1. In configuration mode, go to the following hierarchy level:

   ```
   user@host# edit protocols oam ethernet connectivity-fault-management
     performance-monitoring
   ```

2. Configure the server-side processing option.

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# set delegate-server-processing
   ```

3. Verify the configuration.

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@host# show
   performance-monitoring {
     delegate-server-processing;
   ```

}

See Also

- On-Demand Mode for SLA Measurement on page 732
- *Ethernet Interfaces Feature Guide for Routing Devices*

Related Documentation

- On-Demand Mode for SLA Measurement on page 732
- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768
- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Starting an ETH-DM Session

Using the monitor ethernet delay-measurement Command

After you have configured two MX Series routers to support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), you can initiate a one-way or two-way Ethernet frame delay measurement session from the CFM maintenance association end point (MEP) on one of the routers to the peer MEP on the other router.

To start an ETH-DM session between the specified local MEP and the specified remote MEP, enter the **monitor ethernet delay-measurement** command at operational mode. The syntax of the command is as follows:

monitor ethernet delay-measurement
(one-way | two-way)
maintenance-domain *md-name*
maintenance-association *ma-name*
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<wait *interval-seconds*>
<priority *802.1p value*>
<size>
<no-session-id-tlv>
<xml>

For a one-way frame delay measurement, the command displays a runtime display of the number of 1DM frames sent from the initiator MEP during that ETH-DM session. One-way frame delay and frame delay variation measurements from an ETH-DM session are collected in a CFM database at the router that contains the receiver MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

For a two-way frame delay measurement, the command displays two-way frame delay and frame delay variation values for each round-trip frame exchange during that ETH-DM session, as well as a runtime display of useful summary information about the session: average delay, average delay variation, best-case delay, and worst-case delay. Two-way frame delay and frame delay variation values measurements from an ETH-DM session are collected in a CFM database at the router that contains the initiator MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

NOTE: Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).

For a complete description of the **monitor ethernet delay-measurement** operational command, see the CLI Explorer.

See Also

- monitor ethernet delay-measurement on page 1477

Starting a One-Way ETH-DM Session

To start a one-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement one-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

See Also

- monitor ethernet delay-measurement on page 1477

Starting a Two-Way ETH-DM Session

To start a two-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement two-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

See Also

- monitor ethernet delay-measurement on page 1477

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Starting an ETH-DM Session on page 769
- **monitor ethernet delay-measurement** command
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 771
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 784

Starting a One-Way ETH-DM Session

To start a one-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement one-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

Related Documentation

- monitor ethernet delay-measurement on page 1477

Starting a Two-Way ETH-DM Session

To start a two-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement two-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

Related Documentation

- monitor ethernet delay-measurement on page 1477

Managing ETH-DM Statistics and ETH-DM Frame Counts

Displaying ETH-DM Statistics Only

Purpose Display ETH-DM statistics.

By default, the **show oam ethernet connectivity-fault-management delay-statistics** command displays ETH-DM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain ma1 maintenance-association ma1**

- To display the ETH-DM statistics collected for ETH-DM sessions for the local MEP **201** belonging to MA **ma2** and within MD **md2**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the ETH-DM statistics collected for ETH-DM sessions from local MEPs belonging to MA **ma3** and within MD **md3** to remote MEP **302**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md3 maintenance-association ma3 remote-mep 302**

See Also

- **show oam ethernet connectivity-fault-management delay-statistics on page 2029**

Displaying ETH-DM Statistics and Frame Counts

Purpose Display ETH-DM statistics and ETH-DM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays ETH-DM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma1
```

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP **201** in MA **ma2** and within MD **md2**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md2 maintenance-association ma2 local-mep 201
```

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP in MD **md3** and within MA **ma3** that participates in an ETH-DM session with the remote MEP **302**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain ma3 maintenance-association ma3 remote-mep 302
```

See Also

- show oam ethernet connectivity-fault-management mep-statistics on page 2066

Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain ma1 maintenance-association ma1
```

- To display CFM database information (including ETH-DM frame counts) only for local MEP **201** in MA **ma1** within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md2 maintenance-association ma2 local-mep 201
```

- To display CFM database information (including ETH-DM frame counts) only for remote MEP **302** in MD **md3** within MA **ma3**:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain ma3 maintenance-association ma3 remote-mep 302
```

See Also

- show oam ethernet connectivity-fault-management mep-database on page 2055

Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-DM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces detail**

- To display CFM database information (including ETH-DM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail**

- To display CFM database information (including ETH-DM frame counts) only for MEPs enclosed within CFM maintenance domains (MDs) at level **6**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces level 6 detail**

See Also

- show oam ethernet connectivity-fault-management interfaces on page 2037

Clearing ETH-DM Statistics and Frame Counts

Purpose Clear the ETH-DM statistics and ETH-DM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the ETH-DM statistics and ETH-DM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management statistics**

- To clear the ETH-DM statistics and ETH-DM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

  `user@host>` **clear oam ethernet connectivity-fault-management statistics ge-0/5/9.0**

See Also

- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 784

Related Documentation

- clear oam ethernet connectivity-fault-management statistics on page 1469
- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- show oam ethernet connectivity-fault-management interfaces on page 2037
- show oam ethernet connectivity-fault-management mep-statistics on page 2066
- show oam ethernet connectivity-fault-management mep-database on page 2055

Displaying ETH-DM Statistics Only

Purpose

Display ETH-DM statistics.

By default, the **show oam ethernet connectivity-fault-management delay-statistics** command displays ETH-DM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain ma1 maintenance-association ma1**

- To display the ETH-DM statistics collected for ETH-DM sessions for the local MEP **201** belonging to MA **ma2** and within MD **md2**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the ETH-DM statistics collected for ETH-DM sessions from local MEPs belonging to MA **ma3** and within MD **md3** to remote MEP **302**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md3 maintenance-association ma3 remote-mep 302**

Related Documentation

- show oam ethernet connectivity-fault-management delay-statistics on page 2029

Displaying ETH-DM Statistics and Frame Counts

Purpose

Display ETH-DM statistics and ETH-DM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays ETH-DM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1**

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP **201** in MA **ma2** and within MD **md2**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP in MD **md3** and within MA **ma3** that participates in an ETH-DM session with the remote MEP **302**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Related Documentation

- show oam ethernet connectivity-fault-management mep-statistics on page 2066

Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity

Purpose

Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

`user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma1 maintenance-association ma1**

- To display CFM database information (including ETH-DM frame counts) only for local MEP **201** in MA **ma1** within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display CFM database information (including ETH-DM frame counts) only for remote MEP **302** in MD **md3** within MA **ma3**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Related Documentation

- show oam ethernet connectivity-fault-management mep-database on page 2055

Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-DM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces detail**

- To display CFM database information (including ETH-DM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail**

- To display CFM database information (including ETH-DM frame counts) only for MEPs enclosed within CFM maintenance domains (MDs) at level **6**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces level 6 detail**

Related Documentation

- show oam ethernet connectivity-fault-management interfaces on page 2037

Clearing ETH-DM Statistics and Frame Counts

Purpose Clear the ETH-DM statistics and ETH-DM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the ETH-DM statistics and ETH-DM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management statistics**

- To clear the ETH-DM statistics and ETH-DM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

  `user@host>` **clear oam ethernet connectivity-fault-management statistics ge-0/5/9.0**

Related Documentation

- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 784

Configuring MEP Interfaces

Before you can start an Ethernet frame delay measurement session across an Ethernet service, you must configure two MX Series routers to support ETH-DM.

To configure an Ethernet interface on a MX Series router to support ETH-DM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

   ```
   [edit interfaces]
   interface {
     ethernet-interface-name {
       vlan-tagging;
       unit logical-unit-number {
         vlan-id vlan-id; # Both interfaces on this VLAN
       }
     }
   }
   ```

   Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

   ```
   [edit protocols]
   oam {
   ```



```
ethernet {
  connectivity-fault-management {
    maintenance-domain md-name { # On both routers
      level number;
      maintenance-association ma-name { # On both routers
        continuity-check {
          interval 100ms;
          hold-interval 1;
        }
        mep mep-id { # Attach to VLAN interface
          auto-discovery;
          direction (up | down);
          interface interface-name;
          priority number;
        }
      }
    }
  }
}
}
```

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Ensuring That Distributed ppm Is Not Disabled

By default, the router's period packet management process (**ppm**) runs sessions distributed to the Packet Forwarding Engine in addition to the Routing Engine. This process is responsible for periodic transmission of packets on behalf of its various client processes, such as Bidirectional Forwarding Detection (BFD), and it also receives packets on behalf of client processes.

In addition, **ppm** handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With **ppm** processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run such processes as BFD on the Packet Forwarding Engine.

Distributed ppm Required for ETH-DM

Ethernet frame delay measurement requires that **ppm** remains distributed to the Packet Forwarding Engine. If **ppm** is not distributed to the Packet Forwarding Engines of both routers, ETH-DM PDU frame timestamps and ETH-DM statistics are not valid.

Before you start ETH-DM, you must verify that the following configuration statement is *NOT* present:

```
[edit]
routing-options {
  ppm {
    no-delegate-processing;
  }
}
```

If distributed **ppm** processing is disabled (as shown in the stanza above) on either router, you must re-enable it in order to use the ETH-DM feature.

Procedure to Ensure that Distributed ppm is Not Disabled

To ensure that distributed **ppm** is not disabled on a router:

1. Display the packet processing management (PPM) configuration to determine whether distributed **ppm** is disabled.

   - In the following example, distributed **ppm** is enabled on the router. In this case, you do not need to modify the router configuration:

     ```
     [edit]
     user@host# show routing-options
     ppm;
     ```

   - In the following example, distributed **ppm** is disabled on the router. In this case, you must proceed to Step 2 to modify the router configuration:

     ```
     [edit]
     user@host show routing-options
     ppm {
        no-delegate-processing;
     }
     ```


2. Modify the router configuration to re-enable distributed **ppm** and restart the Ethernet OAM Connectivity Fault Management process *ONLY IF* distributed **ppm** is disabled (as determined in the previous step).

   a. Before continuing, make any necessary preparations for the possible loss of connectivity on the router.

      Restarting the **ethernet-connectivity-fault-management** process has the following effect on your network:

      - All connectivity fault management (CFM) sessions re-establish.
      - All ETH-DM requests on the router terminate.
      - All ETH-DM statistics and frame counts reset to 0.

   b. Modify the router configuration to re-enable distributed **ppm**. For example:

      ```
      [edit]
      user@host# delete routing-options ppm no-delegate-processing
      ```

   c. Commit the updated router configuration. For example:

      ```
      [edit]
      user@host# commit and-quit
      commit complete
      exiting configuration mode
      ```

   d. To restart the Ethernet OAM Connectivity-Fault-Management process, enter the **restart ethernet-connectivity-fault-management <gracefully | immediately | soft>** operational mode command. For example:

      ```
      user@host> restart ethernet-connectivity-fault-management
      Connectivity fault management process started, pid 9893
      ```

Connectivity fault management (CFM) sessions operate in centralized mode over AE interfaces by default. Y.1731 performance monitoring (PM) is supported on centralized CFM sessions over AE interfaces. Also, distribution of CFM session over AE interfaces to line cards is supported from Junos OS Release 13.3. To enable the distribution of CFM sessions and to operate in centralized mode, include the **ppm delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level. The mechanism that enables distribution of CFM sessions over AE interfaces provides the underlying infrastructure to support PM over AE interfaces. In addition, periodic packet management (PPM) handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With PPM processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run performance monitoring processes on the Packet Forwarding Engine.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Enabling the Hardware-Assisted Timestamping Option

By default, Ethernet frame delay measurement uses software for timestamping transmitted and received ETH-DM frames. For Ethernet interfaces, you can optionally use hardware timing to assist in the timestamping of received ETH-DM frames to increase the accuracy of delay measurements.

Enabling hardware-assisted timestamping of received frames can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction.

To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      performance-monitoring {
        hardware-assisted-timestamping;
      }
    }
  }
}
```

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768

Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling

Scaling is the ability of a system to handle increasing amounts of work and to continue to function well. Scaling can refer to increasing capacity and the ability to handle increasing workload, number of subscribers or sessions, hardware components, and so on. Continuity check protocol is used for fault detection within a maintenance association. The maintenance association end points (MEPs) send continuity check messages (CCMs) periodically. The time between the transmissions of CCMs is known as the interval. The receiving MEP maintains a database of all MEPs in the maintenance association.

By default, CCMs are transmitted by the CPU of a line card, such as a Modular Port Concentrator (MPC). If the duration between transmissions of CCMs is low or if the CCMs for a specific line card scale, then we recommend that you delegate transmission of CCMs to the forwarding ASIC (that is, to the hardware) by enabling inline transmission of CCMs. Inline transmission of CCMs is also known as inline keepalives or Inline-KA. Inline transmission enables the system to handle more connectivity fault management (CFM) sessions per line card. By enabling inline transmission of CCMs, you can achieve maximum scaling of CCMs.

To enable inline transmission of CCMs, perform the following steps:

1. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level.

   ```
   [edit]
   user@host# edit protocols oam ethernet connectivity-fault-management
     performance-monitoring
   ```

2. Delegate transmission of CCMs to hardware by enabling hardware-assisted keepalives.

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# set hardware-assisted-keepalives enable
   ```

   NOTE: Inline transmission of CCMs is not enabled when there is a CFM session already established. To enable inline transmission, you must first deactivate the CFM session using the **deactivate** command and then reactivate the CFM session using the **activate** command.

To disable inline transmission, use the hardware-assisted-keepalives **disable** statement. After disabling inline transmission, you must reboot the router for the changes to take effect.

Related Documentation

- Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades on page 800

Enabling Inline Mode Of Performance Monitoring To Achieve Maximum Scaling

Performance monitoring is useful for studying the traffic pattern in a network over a period of time. It helps to identify network problems before you are impacted by network defects.

By default, performance monitoring packets are handled by the CPU of a line-card, such as Modular Port Concentrator (MPC). Enabling inline mode of performance monitoring delegates the processing of the protocol data units (PDUs) to the forwarding ASIC (that is, to the hardware). By enabling inline mode of performance monitoring, the load on the CPU of the line-card is reduced and you can configure an increased number of performance monitoring sessions and achieve maximum scaling for service OAM performance monitoring sessions. On MX Series routers, you can configure inline mode of performance monitoring only if the network services mode on the router is configured to **enhanced-ip** and enhanced connectivity fault management (**enhanced-cfm-mode**)) is configured.

By enabling inline mode of performance monitoring, you can achieve maximum scaling for performance monitoring sessions. To achieve maximum scaling for performance monitoring sessions, you must enable scaling of continuity check messages (CCMs) sessions. To enable scaling of CCM sessions, enable inline transmission of continuity check messages. For more information on inline transmission of continuity check messages, see "Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling" on page 795. To view the supported scaling values for CCM and PM, see "Supported Inline CCM and Inline PM Scaling Values" on page 798.

Inline mode of performance monitoring is supported only for proactive mode of frame delay measurement (Two-way Delay Measurements) and synthetic loss measurements (SLM)sessions. Performance monitoring functions configured using the iterator profile (CFM)are referred to as proactive performance monitoring. Inline mode of performance monitoring for frame loss measurement using service frames (LM) is not supported.

NOTE: **MPC3E (MX-MPC3E-3D) and MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE) do not support inline mode of performance monitoring. User-defined Data TLV is not supported if you have configured inline mode of performance monitoring. Also, only 12 history records per PM sessions are supported.**

We recommend that you enable inline mode of performance monitoring before you configure the performance monitoring sessions as the change may interfere with the existing performance monitoring sessions.

To enable inline mode of performance monitoring, perform the following steps:

1. In configuration mode, go to the **[edit chassis]** hierarchy level and configure the network services mode of the router. The network service mode of the router must be configured as **enhanced ip** to enable enhanced connectivity fault management (CFM) mode.

NOTE: If the network services mode is not **enhanced-ip**, and you have enabled enhanced CFM, the following warning message is displayed:
[edit protocols oam ethernet]
'connectivity-fault-management'
enhanced ip is not effective please configure enhanced ip and give router reboot

```
[edit chassis]
user@host# set network-services enhanced-ip
```

2. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level and enable enhanced connectivity fault management mode by using the **enhanced-cfm-mode** option.

```
[edit]
user@host# set protocols oam ethernet connectivity-fault-management
  enhanced-cfm-mode
```

3. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level. Configure the enhanced iterator profile by using the **enhanced-sla-iterator** option and specify the measurement interval by using the **measurement-interval** option.

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management
  performance-monitoring enhanced-sla-iterator measurement-interval value
```

4. Enable inline performance monitoring.

NOTE: You can enable inline mode of performance monitoring for both the originator and the responder of the service OAM performance monitoring sessions by using the **hardware-assisted-pm** command.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-pm
```

5. (Optional) Enable inline transmission of CCMs to enable better scaling if inline transmission of CCMs is not automatically enabled.

NOTE: You can achieve better scaling if both inline performance monitoring and inline transmission of CCMs is enabled.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-keepalives enable
```

6. Commit the configuration.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# commit
```

Related Documentation

- Enabling Enhanced Connectivity Fault Management Mode on page 666
- Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling on page 795
- *Network Services Mode Overview*
- *hardware-assisted-pm*
- Supported Inline CCM and Inline PM Scaling Values on page 798

Supported Inline CCM and Inline PM Scaling Values

This topic lists the scaling values for inline mode of performance monitoring and inline transmission of continuity check messages. The scaling values are based on the different cycle-time interval values. Each table lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure inline CCM, enhanced CFM, and enhanced PM by using the **hardware-assisted-keepalives**, **enhanced-cfm-mode**, and **hardware-assisted-pm** options.

NOTE: The scaling values do not consider the load from other protocols in the system and so the actual realized scaling values for line card and chassis vary depending on other protocol configurations and scaling in the system. We recommend that you configure DDoS for CFM. Limit the number of CFM packets, that are sent to the CPU of the line card, to 3000. Limiting the number of packets safeguards the CPU from scaled CFM configurations of various CFM protocol events.

Table 75 on page 798 lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure both the CCM interval and the PM interval as 1 second.

Table 75: Scaling Values for CFM and PM (CCM Interval: 1 sec and PM Interval: 1 sec)

CFM Line Card Scale	PM Line Card Scale	CFM Chassis Scale	PM Chassis Scale
4000	4500	16000	16000
6000	3750	16000	16000
7000	3375	16000	16000
8000	3000	16000	16000

Table 76 on page 799 lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure the CCM interval as 1 second and the PM interval as 100 milliseconds.

Table 76: Scaling Values for CFM and PM (CCM Interval: 1 sec and PM interval: 100 ms)

CFM Line Card Scale	PM Line Card Scale	CFM Chassis Scale	PM Chassis Scale
4000	450	12000	4000
6000	375	12000	4000
7000	337	12000	4000
8000	300	12000	4000

Table 77 on page 799 lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure the CCM interval as 100 milliseconds and the PM interval as 1 second.

Table 77: Scaling Values for CFM and PM (CCM Interval: 100 ms and PM interval: 1 sec)

CFM Line Card Scale	PM Line Card Scale	CFM Chassis Scale	PM Chassis Scale
4000	3000	8000	6000
3000	3750	8000	6000
2000	4500	8000	6000
1000	4500	8000	6000

Table 78 on page 799 lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure both the CCM interval and the PM interval as 100 milliseconds.

Table 78: Scaling Values for CFM and PM (CCM Interval: 100 ms and PM interval: 100 ms)

CFM Line Card Scale	PM Line Card Scale	CFM Chassis Scale	PM Chassis Scale
4000	300	8000	3000
3000	375	8000	3000
2000	450	8000	3000
1000	450	8000	3000

Related Documentation

- *hardware-assisted-pm*

- Enabling Inline Mode Of Performance Monitoring To Achieve Maximum Scaling on page 796

Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades

Starting in Release 17.1, Junos OS connectivity fault management (CFM), during a unified in-service software upgrade (ISSU), works when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit continuity check message (CCM) PDUs during the unified ISSU. Connectivity fault management continues to operate.

This feature requires the following conditions be met:

- Packet Forwarding Engine keepalives must be enabled to provide inline transmission of CCMs. The feature does not work when the CCMs are transmitted by the CPU of a line card, which is the default transmission method.
- The interval between CCMs must be 1 second.

CFM interoperability during a unified ISSU is supported on the following MPCs: MPC1, MPC2, MPC2-NG, MPC3-NG, MPC5, and MPC6.

To enable CFM interoperability with third-party devices across a unified ISSU:

1. Enable inline keepalives.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-keepalives enable
```

2. Set the CCM interval to 1 second.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name continuity-check]
user@host# set interval 1s
```

Release History Table

Release	Description
17.1	Starting in Release 17.1, Junos OS connectivity fault management (CFM), during a unified in-service software upgrade (ISSU), works when the peer device is not a Juniper Networks router.

Related Documentation

- Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling on page 795
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612

Using the monitor ethernet delay-measurement Command

After you have configured two MX Series routers to support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), you can initiate a one-way or two-way Ethernet frame delay measurement session from the CFM maintenance association end point (MEP) on one of the routers to the peer MEP on the other router.

To start an ETH-DM session between the specified local MEP and the specified remote MEP, enter the **monitor ethernet delay-measurement** command at operational mode. The syntax of the command is as follows:

monitor ethernet delay-measurement
(one-way | two-way)
maintenance-domain *md-name*
maintenance-association *ma-name*
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<wait *interval-seconds*>
<priority *802.1p value*>
<size>
<no-session-id-tlv>
<xml>

For a one-way frame delay measurement, the command displays a runtime display of the number of 1DM frames sent from the initiator MEP during that ETH-DM session. One-way frame delay and frame delay variation measurements from an ETH-DM session are collected in a CFM database at the router that contains the receiver MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

For a two-way frame delay measurement, the command displays two-way frame delay and frame delay variation values for each round-trip frame exchange during that ETH-DM session, as well as a runtime display of useful summary information about the session: average delay, average delay variation, best-case delay, and worst-case delay. Two-way frame delay and frame delay variation values measurements from an ETH-DM session are collected in a CFM database at the router that contains the initiator MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

NOTE: Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).

For a complete description of the **monitor ethernet delay-measurement** operational command, see the CLI Explorer.

Related Documentation

- monitor ethernet delay-measurement on page 1477

Managing ETH-LM Statistics

Displaying ETH-LM Statistics

Purpose Display the ETH-LM statistics.

By default, the **show oam ethernet connectivity-fault-management loss-statistics maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*** command displays ETH-LM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

The following list consists of the CFM-related operational mode commands that have been enhanced to display ETH-LM statistics:

- The **show oam ethernet connectivity-fault-management interfaces detail** command is enhanced to display ETH-DM and ETH-LM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).
- The **show oam ethernet connectivity-fault-management mep-statistics** command is enhanced to display ETH-DM and ETH-LM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).
- The **show oam ethernet connectivity-fault-management mep-database** command is enhanced to display ETH-DM and ETH-LM frame counters for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-LM statistics for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management loss-statistics**

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1**

- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1**

- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

`user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain md1 maintenance-association ma1**

See Also

- clear oam ethernet connectivity-fault-management loss-statistics on page 1467
- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- show oam ethernet connectivity-fault-management interfaces on page 2037 (**detail** | **extensive**)
- show oam ethernet connectivity-fault-management mep-statistics on page 2066
- show oam ethernet connectivity-fault-management mep-database on page 2055
- show oam ethernet connectivity-fault-management loss-statistics on page 2051
- *Ethernet Interfaces Feature Guide for Routing Devices*

Clearing ETH-LM Statistics

Purpose Clear the ETH-LM statistics.

By default, statistics are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the ETH-LM statistics for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management loss-statistics**

See Also

- clear oam ethernet connectivity-fault-management loss-statistics on page 1467
- *Ethernet Interfaces Feature Guide for Routing Devices*

Related Documentation

- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 784

Managing Continuity Measurement Statistics

- Displaying Continuity Measurement Statistics on page 803
- Clearing Continuity Measurement Statistics on page 804

Displaying Continuity Measurement Statistics

Purpose Display continuity measurement.

The **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1** command is enhanced to display continuity measurement statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1**

See Also

- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- *Ethernet Interfaces Feature Guide for Routing Devices*

Clearing Continuity Measurement Statistics

Purpose Clear the continuity measurement statistics

By default, statistics are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the continuity measurement statistics for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management continuity-measurement maintenance-domain *md-name* maintenance-association *ma-name* local-mep *local-mep-id* remote-mep *remote-mep-id***

See Also

- clear oam ethernet connectivity-fault-management continuity-measurement on page 1465
- *Ethernet Interfaces Feature Guide for Routing Devices*

Related Documentation

- clear oam ethernet connectivity-fault-management continuity-measurement on page 1465
- show oam ethernet connectivity-fault-management delay-statistics on page 2029

Configuring the Failure Notification Protocol

This topic describes how to configure the Ethernet Operations, Administration, and Maintenance (OAM) Failure Notification Protocol (FNP) on MX Series routers. The FNP detects link failures in a Carrier Ethernet network and broadcasts FNP messages when a failure occurs to all nodes affected by the link failure. To configure FNP functionality, include the **fnp** statement at the **[edit protocols oam ethernet]** hierarchy level:


```
[edit protocols oam]
ethernet {
   fnp {
      interval <100ms | 1s | 10s | 1m | 10m>;
      loss-threshold number
      interface interface name {
         domain-id domain-id
      }
   }
}
```

The **interval** statement specifies the time between the transmission of FNP messages. You can specify 10 minutes (10m), 1 minute (1m), 10 seconds (10s), 1 second (1s), and 100 milliseconds (100ms). The **loss-threshold** statement specifies how many FNP messages can be lost before the FNP message is considered aged out and flushed. You must include the **interface *interface-name*** statement with the **domain-id *domain-id*** statement. The **domain-id** statement specifies a domain ID for the route. FNP messages can be received and processed on MX Series routers, but generating FNP messages is not supported.

The **show oam ethernet fnp interface**, **show oam ethernet fnp status**, and **show oam ethernet fnp messages** operational commands display the configured information.

FNP can be enabled only on logical interfaces that are part of a VPLS routing instance, and none of the logical interfaces in the VPLS routing instance should have CCM configured. FNP can be enabled on only one logical interface per physical interface.

Related Documentation

- **connectivity-fault-management on page 1113**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

Ethernet alarm indication signal (ETH-AIS) function enables a service provider deploying an Ethernet service to determine whether a connectivity fault exists at the provider's domain level or at a level below. When the fault occurs at the provider's domain level, the service provider addresses the fault, and when the fault occurs at a level below, the provider can either ignore the fault or contact the relevant authorities to address the fault.

The following sections explain ETH-AIS, few use cases which determine when to generate and propagate ETH-AIS packets, and associated terms in detail:

Understanding ETH-AIS in a Maintenance Domain

ITU-T developed Y.1731 as a recommendation for Operation, Administration, and Maintenance (OAM) functions and mechanisms for Ethernet-based networks, including

OAM functions such as ETH-AIS, Ethernet locked signal (ETH-LCK), Ethernet test signal (ETH-Test), Ethernet automatic protection switching (ETH-APS), Ethernet maintenance communication channel (ETH-MCC), Ethernet experimental OAM (ETH-EXP), Ethernet vendor-specific OAM (ETH-VSP), and performance monitoring. For information about maintenance domain and related terms, see "Terms Defined" on page 808.

According to the Y.1731 standards, a server MEP is a combined function of the server layer termination function and the server Ethernet services layer adaptation function. The server MEP notifies the Ethernet services (ETH) layer MEPs when it detects a failure. The server layer termination function then runs the OAM mechanisms specific to the server layer and the alarms are suppressed at the server layer by ETH-AIS.

Note that ETH-AIS is not applicable to Spanning Tree Protocol (STP) networks.

ETH-AIS enables you to suppress alarms when a fault condition is detected. Using ETH-AIS, a service provider can differentiate between faults at different levels.

ETH-AIS provides many advantages that include:

- Service providers need not raise alarms if there are lower-level failures.
- Service providers can provide a refund to their subscribers or avail a refund from their Internet provider based on service unavailability.

MX Series routers support ITU-T Y.1731 ETH-AIS to provide fault management for service providers who provide carrier Ethernet services using IEEE 802.1ag standard.

NOTE: MX Series Virtual Chassis does not support Ethernet alarm indication signal (ETH-AIS).

Fault Detection in a Maintenance Domain

In the scenario depicted in Figure 1 on page xyz, you have a service provider level and a customer level. Two service providers—*Operator-1* and *Operator-2*—are considered for illustration purposes. Assume that a fault occurs in Operator-1 maintenance domain-level that has MEP-A and MEP-B at its maintenance domain-level boundaries. To notify the faults to a network management system and to avoid notification of alarms from the customer level for the same fault, MEP-A and MEP-B transmit an alarm indication signal (AIS) on opposite directions, thereby signaling the higher levels and the Operator-2 network about the fault, so that the alarms are suppressed.

Signaling is achieved through transmission and propagation of AIS protocol data units (PDUs). You must enable AIS explicitly on all the MEPs at the service provider level. A MEP that is configured to issue frames with ETH-AIS information is generally at the server layer and continues to transmit periodic frames with ETH-AIS information until the defect condition is cleared. When a client MEP receives the ETH-AIS frames, it suppresses loss-of-continuity alarms associated with its peer MEPs.

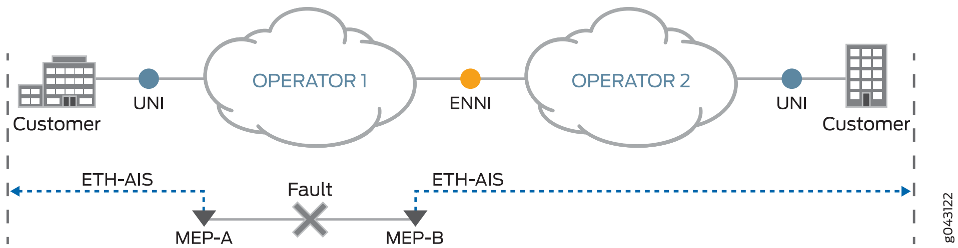

Note that in the absence of AIS, a client MEP resumes generating loss-of-continuity alarms when it detects the loss-of-continuity defect conditions from its server layer.

For point-to-point Ethernet services layer connectivity, a MEP has only one peer MEP. Therefore, there is no ambiguity regarding the peer MEP for which the MEP should suppress alarms when it receives the ETH-AIS information.

For multipoint Ethernet services layer connectivity, a MEP that receives ETH-AIS information cannot determine the exact MEP that encountered the fault and, therefore, cannot isolate the exact peer MEP to suppress the alarms. To avoid this scenario, Y.1731 recommends suppressing alarms for all peer MEPs in the same domain level irrespective of connectivity status in a multipoint Ethernet services layer connectivity setup.

Table 79 on page 807 lists the operational mode commands that you can use in a maintenance domain to check the various parameters pertaining to a MEP.

Table 79: Operational Mode Commands

To Check	Operational Mode Commands
Whether the AIS configuration is configured correctly on a CFM MEP.	**show protocols oam ethernet connectivity-fault-management action-profile**
Statistics of AIS frames.	**show oam ethernet connectivity-fault-management interfaces detail** **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *md-name* maintenance-association *ma-name* remote-mep *mep-id* local-mep *mep-id***
Whether any event has occurred that triggered AIS.	**show oam ethernet connectivity-fault-management mep-database maintenance-domain *md-name* maintenance-association *ma-name* remote-mep *mep-id* local-mep *mep-id***
Status of CFM sessions for faults that trigger AIS on the MEP.	**show oam ethernet connectivity-fault-management interfaces detail**

Terms Defined

- AIS transmission—A MEP upon detecting a defect condition transmits AIS frames in a direction opposite to its peer MEPs. The periodicity of AIS frames transmission is on the basis of the AIS transmission period. An AIS transmission period of 1 second is recommended. The first AIS frame must always be transmitted immediately following the detection of a defect condition.
- AIS reception—Upon receiving an AIS frame, a MEP examines it to ensure that the frame's maintenance domain level is the same as its own maintenance domain level. The *period* field in the frame indicates the period at which the AIS frames can be expected. When a MEP receives an AIS frame, it detects the defect condition. After detection, when no AIS frames are received within an interval of 3.5 times—the AIS transmission period indicated in the AIS frames received—the MEP clears the AIS defect condition. When the AIS condition is cleared and defects still exist, then the MEPs continue to report alarms.
- AIS PDU format—The fields of the AIS PDU format are:

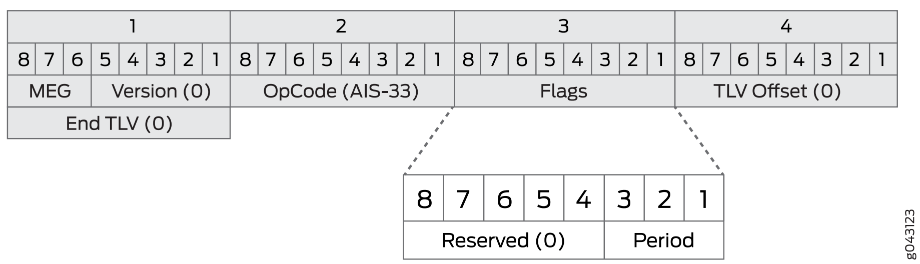

a. MEG Level—Also called the maintenance domain level, it is a 3-bit field that is used to carry the maintenance domain level of the client MEG.

b. Version—Value is always 0.

c. OpCode—Value for this PDU type is AIS (33).

d. Flags—The first five bits are reserved and are set to 0. The 3-bit information element carried in the three least significant bits are referred to as the period that contains the value of AIS transmission periodicity as illustrated in Table 80 on page 808:

Table 80: AIS Transmission Periodicity

Flags [3:1]	Period Value	Comments
000-011	Invalid value	Invalid value for AIS
100	1s	1 frame per second
101	Invalid value	Invalid value for AIS
110	1 min	1 frame per minute
111	Invalid value	Invalid value for AIS

e. TLV offset—Set to 0.

f. End TLV—All-zeroes octet value.

- Server layer and client layer—These layers are part of the ITU-T Recommendation G.805 transport network functional model. This model is based on the concept of layering within a transport network. A transport network is divided into several independent transport layer networks that have a client-server association between adjacent layer networks.
- Maintenance domain—To enable connectivity fault management (CFM) on an Ethernet interface, maintenance domains, maintenance associations, and maintenance end points (MEPs) are created and configured in a network. You can configure up to eight maintenance domain levels in a network. Each maintenance domain level is a part of the network where the connectivity issues can be monitored and corrected. Provider domain and customer domain are some examples for maintenance domains. Each maintenance domain has a maintenance association. Each maintenance association includes MEPs and maintenance intermediate points (MIPs) in that domain. The MEPs are located at the boundary of the domain and the MIPs are located within the domain. MEPs generate and transmit continuity check messages (CCMs) at configured intervals to the entire maintenance association to check the connectivity in the network.
- Ethernet services (ETH) layer—A layer in the metro Ethernet network model, where this layer is responsible for the OAM services that are required to support the Ethernet services in the network.

Related Documentation

- Configuring ETH-AIS on a CFM MEP on page 811

Ethernet Alarm Indication Signal Overview

ACX Series routers support ITU-T Y.1731 Ethernet Alarm Indication Signal function (ETH-AIS) to provide fault management for service providers. ETH-AIS enables you to suppress alarms when a fault condition is detected. Using ETH-AIS, an administrator can differentiate between faults at customer level or faults at provider level.

The advantage of ETH-AIS is:

- Customers need not raise alarms due to lower level failures.
- Customers can get refund based on service unavailability.

When a fault condition is detected, a maintenance end point (MEP) generates ETH-AIS packets to the configured client levels for a specified duration until the fault condition is cleared. Any MEP configured to generate ETH-AIS packets signals to a level higher than its own. A MEP receiving ETH-AIS recognizes that the fault is at a lower level and then suppresses alarms at current level.

ACX Series routers support ETH-AIS PDU generation for server MEPs based on the following defect conditions:

- Loss of connectivity (physical link loss detection)

- Layer 2 circuit or Layer 2 VPN down

Alarm indication signaling is done through the transmission and propagation of ETH-AIS PDUs. ETH-AIS should be enabled on MEPs. A MEP which is configured to issue packets with ETH-AIS information is generally of server layer and continues to transmit periodic packets with ETH-AIS information until the defect condition is cleared. CFM MEPs, upon receiving ETH-AIS PDUs, suppresses loss of continuity alarms associated with its peer MEPs. A MEP resumes loss of continuity alarm generation upon detecting loss of continuity defect conditions in the absence of an ETH-AIS condition.

For point-to-point Ethernet connectivity, a MEP has only a single peer MEP. Therefore, a MEP suppress alarms on its peer MEP when it receives the ETH-AIS information.

For multi-point Ethernet connectivity, a MEP which receives ETH-AIS information cannot determine the exact MEP encountered a fault condition and therefore it will not be able to isolate the exact peer MEP for alarm suppression. ITU-T Y.1731 recommends suppressing alarms for all peer MEPs irrespective of the connectivity status.

AIS transmission—A MEP upon detecting a defect condition transmits ETH-AIS PDUs in a direction opposite to its peer MEPs. The transmission of ETH-AIS PDUs is based on a configured ETH-AIS transmission period. An ETH-AIS transmission period of 1 second is recommended. The first ETH-AIS PDU must be transmitted immediately following the detection of a defect condition.

AIS reception—A MEP upon receiving ETH-AIS PDUs examines it to ensure that its maintenance domain (MD) level corresponds to the same MD level. Upon receiving an ETH-AIS PDU, the MEP detects a defect condition. Following the detection of a defect condition, if there are no ETH-AIS PDUs received within an interval of 3.5 times the ETH-AIS transmission period indicated in the ETH-AIS PDUs received earlier, the MEP clears the defect condition. After the fault condition is cleared, MEPs continue to report alarms.

NOTE: ACX Series routers do not support ITU-T Y.1731 ETH-AIS for layer 2 services (bridging).

The following are the limitations for server MEP

- Triggering of ETH-AIS messages over services (Layer 2 circuit and Layer 2 VPN) by the link-loss server MEP is done on a best-effort manner. This is because the transmission of ETH-AIS messages is independent of the service status and there is no guarantee for delivering the ETH-AIS messages before service goes down.
- Pseudowire protection with CFM-MEP session is not monitored by the server-MEP because an entity to monitor pseudowire protection already exists for the service (Layer 2 circuit and Layer 2 VPN).

Related Documentation

- Configuring Alarm Indication Signal on ACX Series Routers on page 815

Configuring ETH-AIS on a CFM MEP

MX Series routers support ITU-T Y.1731 Ethernet alarm indication signal (ETH-AIS) function to provide fault management for service providers. ETH-AIS enables the service provider to suppress alarms when a fault condition is detected.

The following points are to be noted when ETH-AIS is configured in a maintenance domain:

- Transmitting or receiving of AIS on a MEP does not override the **lowest-priority-defect** statement configured at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level. Therefore, alarms are reported according to the defect priority configured.
- Alarms are reported even when the higher domain levels exchange CCMs at a faster rate than the lower domain levels.
- Maintenance association intermediate point (MIP) is transparent to ETH-AIS frames—that is, the MIPs do not perform any action in response to ETH-AIS frames.
- When the service provider requests the MEP to generate an AIS for a lower level or for the same level, the request is rejected.
- AIS generation is stopped when the MEP clears the remote MEP within the maintenance association.
- When the **auto-discovery** statement is enabled for a MEP, the remote MEP information is cleared after the configured hold interval expires.

The following tasks explain how to enable ETH-AIS in a maintenance domain, configure an action to be taken when a defect is detected, and to attach the action profile to a CFM MEP:

1. Configuring an Action Profile on page 811
2. Configuring an Action to Be Taken When an AIS Alarm Is Detected on page 812
3. Attaching the Action Profile to a CFM MEP on page 813

Configuring an Action Profile

To configure an action profile for ETH-AIS:

1. Go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

   ```
   [edit]
   user@host# edit protocols oam ethernet connectivity-fault-management
   ```

2. Configure an action profile to use when one or more remote MEPs are down.

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

3. Configure an event that needs to be monitored.

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

4. Configure the defect condition that generates an alarm indication signal.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event]
   user@host# **edit ais-trigger-condition**

5. Configure the **adjacency-loss** statement to inform the operator when the physical connectivity is lost between the peer MEPs.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]
   user@host# **set adjacency-loss**

6. Configure the **all-defects** statement to inform the operator that all possible defects must be considered to raise the alarm indication signal.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]
   user@host# **set all-defects**

7. Configure the **cross-connect-ccm** statement to inform the operator when cross-connect continuity check messages (CCMs) are received by the MEP and to raise an alarm indication signal in response.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]
   user@host# **set cross-connect-ccm**

8. Configure the **erroneous-ccm** statement to inform the operator when CCMs with unexpected MEP ID or maintenance domain level are received by the MEP and an AIS alarm is raised in response.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]
   user@host# **set erroneous-ccm**

9. Configure the **receive-ais** statement to inform the operator that an AIS message has been received from the peer MEP in its own maintenance level.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]
   user@host# **set receive-ais**

Configuring an Action to Be Taken When an AIS Alarm Is Detected

Configure an action to be taken when an AIS alarm is detected.

1. Go to the **[edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action]** hierarchy level.

[edit]
user@host# **edit protocols oam ethernet connectivity-fault-management action-profile** *action-profile-name* **action**

2. Configure the **log-and-generate-ais** statement to log the event that generated the AIS message.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action]
   user@host# **edit log-and-generate-ais**

3. Configure the interval between AIS messages that are to be received by the MEP as 1 minute or 1 second.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action log-and-generate-ais]
   user@host# **set interval (1m | 1s)**

4. Configure the server maintenance domain level range of the MEP from 1 through 7.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action log-and-generate-ais]
   user@host# **set level** ***level***

5. Configure the 802.1p priority of the AIS packet from 1 through 7.

   [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action log-and-generate-ais]
   user@host# **set priority** ***level***

Attaching the Action Profile to a CFM MEP

After configuring an event and an action to be monitored in an action profile, you must attach the action profile to a CFM MEP.

1. Go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

   [edit]
   user@host# **edit protocols oam ethernet connectivity-fault-management**

2. Configure the maintenance domain with a name.

   [edit protocols oam ethernet connectivity-fault-management]
   user@host# **edit maintenance-domain** ***md-name***

3. Configure the maintenance domain with a client maintenance entity group (MEG) level or maintenance association level—the level which the client layer maintenance association intermediate point (MIPs) and the MEPs exist—from 0 through 7.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name*]
   user@host# **edit level** ***level***

NOTE: You cannot configure a maintenance domain level that is lower than or equal to the maintenance association level that it is associated with.

4. Configure the maintenance association.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name*]
   user@host# **edit maintenance-association** ***ma-name***

5. Configure the continuity check that is performed on all the MEPs in a domain level by sending CCMs with an interval between two CCMs—100 miiliseconds, 10 milliseconds, 1 second, 10 seconds, 1 minute, or 10 minutes—and the number of CCMs that are to be lost before marking a MEP as down.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]
   user@host# **set continuity-check interval (100ms | 10m | 10ms | 1m | 1s)**
   user@host# **set continuity-check loss-threshold** ***value***

6. Configure the MEP with an identifier from 1 through 8192.

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

7. Attach the configured action profile to the MEP.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]
   user@host# **set action-profile** ***action-profile-name***

8. Configure the interface of the MEP over which the CCMs are transmitted.

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

9. Configure the direction for the CCMs to travel to the next MEP as up or down.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]
   user@host# **set direction (down | up)**

10. Configure the 802.1p priority for the CCMs and link-trace packet from 0 through 7.

    [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]
    user@host# **set priority** ***priority-value***

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805

Configuring Alarm Indication Signal on ACX Series Routers

ACX Series routers support ITU-T Y.1731 Ethernet Alarm Indication Signal function (ETH-AIS) to provide fault management for service providers. ETH-AIS enables you to suppress alarms when a fault condition is detected.

To support ETH-AIS transmission, the following configuration information is required by a CFM MEP:

- Client Maintenance Entity Group level—Maintenance Entity Group (MEG) level at which the immediate client layer Maintenance Domain Intermediate Points (MIPs) and Maintenance Association End Points (MEPs) exist.
- ETH-AIS transmission period—Determines the ETH-AIS PDU transmission interval.
- Priority—Determines the priority of packets with ETH-AIS information. This is optional.

To configure ETH-AIS in CFM MEP, you need to:

- Configure an action profile with ETH-AIS action
- Attach the action profile to the CFM MEP

To configure an action profile with ETH-AIS action, include the following statements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile action-profile-name {
  event {
    adjacency-loss;
    all-defects;
    cross-connect-ccm;
    errored-ccm;
    receive-ais;
  }
  action {
    log-and-generate-ais {
      level [1-7];
      interval 1s | 1m ;
      priority [0-7];
    }
    log-ais;
  }
}
```

To attach an action profile to a CFM MEP, include the following statements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```
maintenance-domain maintenenace-domain-name {
level level-number;
  maintenance-association maintenance-domain-name {
    continuity-check {
      interval 1s;
```

```
            loss-threshold 3;
          }
          mep mep-id {
            interface interface-name;
            direction up | down;
            priority priority-value;
            action-profile action-profile-name;
          }
        }
      }
```

NOTE: You cannot configure a maintenance domain level that is lower than or equal to the level that it is associated with.

To support ETH-AIS transmission, the following configuration information required by a server MEP:

- Server MEP definition—Defines the association of server MEP identifier to the server layer.
  - For Layer 2 circuit and Layer 2 VPN, the logical interface connected to a customer network (UNI) would be the identifier for the server layer that needs to be monitored by the server MEP.
  - For physical link loss detection, the physical interface under Ethernet protocol would be the identifier for the server layer that needs to be monitored by the server MEP.
- Association of server MEP defect—Defines the association of server MEP defects to ETH-AIS action.
- Association action profile and server MEP—Defines the binding of server MEP and action profile.

To configure ETH-AIS in server MEP, you need to:

- Create an action profile with ETH-AIS action for server MEP defects.
- Attach the action profile to a server MEP

To create an action profile, include the following statements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile action-profile-name {
  event {
    server-mep-defects {
      link-loss-defect;
      l2circuit-defect;
      l2vpn-defect;
    }
  }
  action {
    log-and-generate-ais {
      level 1...n;
      interval 1 second | 1 minute;
```



```
            priority dot1p [range 0-7];
        }
    }
}
```

To attach an action profile to a server MEP, include the following statement at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
server-mep mep-identifier {
    protocol l2circuit | l2vpn | ethernet {
        interface interface-name;
    }
    action-profile action-profile-name;
}
```

Related Documentation

- Ethernet Alarm Indication Signal Overview on page 809

Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces

This example uses two MX Series routers: **MX-1** and **MX-2**. The configuration creates a CFM down MEP session on a VLAN-tagged logical interface connecting the two (**ge-5/2/9** on Router **MX-1** and **ge-0/2/5** on Router **MX-2**).

NOTE: These are not complete router configurations.

Configuration on Router **MX-1**:

```
[edit]
interfaces {
    ge-5/2/9 {
        vlan-tagging;
        unit 0 {
            vlan-id 512;
        }
    }
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                traceoptions {
                    file eoam_cfm.log size 1g files 2 world-readable;
                    flag all;
                }
                linktrace {
                    path-database-size 255;
                    age 10s;
                }
                maintenance-domain md6 {
                    level 6;
```

```
                maintenance-association ma6 {
                    continuity-check {
                        interval 100ms;
                        hold-interval 1;
                    }
                    mep 201 {
                        interface ge-5/2/9.0;
                        direction down;
                        auto-discovery;
                    }
                }
            }
        }
    }
}
```

Configuration on Router **MX-2**:

```
[edit]
interfaces {
    ge-0/2/5 {
        vlan-tagging;
        unit 0 {
            vlan-id 512;
        }
    }
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                traceoptions {
                    file eoam_cfm.log size 1g files 2 world-readable;
                    flag all;
                }
                linktrace {
                    path-database-size 255;
                    age 10s;
                }
                maintenance-domain md6 {
                    level 6;
                    maintenance-association ma6 {
                        continuity-check {
                            interval 100ms;
                            hold-interval 1;
                        }
                        mep 101 {
                            interface ge-0/2/5.0;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
            }
        }
```



```
  }
}
```

From Router **MX-2**, start a one-way delay measurement to Router **MX-1**.

```
user@MX-2> monitor ethernet delay-measurement one-way mep 201 maintenance-domain md6
maintenance-association ma6 count 10
One-way ETH-DM request to 00:90:69:0a:43:94, Interface ge-0/2/5.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

The counters are displayed as part of the local MEP database on Router **MX-2**.

```
user@MX-2> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-domain ma6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 1590
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 10
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      201     00:90:69:0a:43:94       ok    ge-0/2/5.0
```

The remote MEP database statistics are available on Router **MX-1**.

```
user@MX-1> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
```

```
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 201, Direction: down, MAC address: 00:90:69:0a:43:94
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 1572
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 0
    Valid 1DMs received                           : 10
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      101     00:90:69:0a:48:57       ok    ge-5/2/9.0
```

The remote Router **MX-1** should also collect the delay statistics (up to 100 per session) for display with **mep-statistics** or **delay-statistics**.

`user@MX-1>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md6**

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
  CCMs sent                                     : 3240
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           : 0
  1DMs sent                                     : 0
  Valid 1DMs received                           : 10
  Invalid 1DMs received                         : 0
  DMMs sent                                     : 0
  DMRs sent                                     : 0
  Valid DMRs received                           : 0
  Invalid DMRs received                         : 0
```



```
  Remote MEP identifier: 101
  Remote MAC address: 00:90:69:0a:48:57
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
             (usec)         (usec)
      1        370
      2        357
      3        344
      4        332
      5        319
      6        306
      7        294
      8        281
      9        269
      10       255
    Average one-way delay          : 312 usec
    Average one-way delay variation: 11 usec
    Best case one-way delay        : 255 usec
    Worst case one-way delay       : 370 usec
```

user@MX-1> **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md6**

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1

  Remote MAC address: 00:90:69:0a:48:57
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
             (usec)         (usec)
      1        370
      2        357
      3        344
      4        332
      5        319
      6        306
      7        294
      8        281
      9        269
      10       255
    Average one-way delay          : 312 usec
    Average one-way delay variation: 11 usec
    Best case one-way delay        : 255 usec
```

NOTE: When two systems are close to each other, their one-way delay values are very high compared to their two-way delay values. This is because one-way delay measurement requires the timing for the two systems to be synchronized at a very granular level and MX Series routers do not support this granular synchronization. However, two-way delay measurement does not require synchronized timing, making two-way delay measurements more accurate.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 723

Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces

This example uses two MX Series routers: **MX-1** and **MX-2**. The configuration creates a CFM down MEP session on a VLAN-tagged logical interface connecting the two (**ge-5/2/9** on Router **MX-1** and **ge-0/2/5** on Router **MX-2**).

NOTE: These are not complete router configurations.

Configuration on Router **MX-1**:

```
[edit]
interfaces {
  ge-5/2/9 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        linktrace {
          path-database-size 255;
          age 10s;
        }
        maintenance-domain md6 {
          level 6;
          maintenance-association ma6 {
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 201 {
              interface ge-5/2/9.0;
              direction down;
              auto-discovery;
```



```
                }
              }
            }
          }
        }
      }
    }
```

Configuration on Router **MX-2**:

```
[edit]
interfaces {
  ge-0/2/5 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        linktrace {
          path-database-size 255;
          age 10s;
        }
        maintenance-domain md6 {
          level 6;
          maintenance-association ma6 {
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 101 {
              interface ge-0/2/5.0;
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

From Router **MX-1**, start a two-way delay measurement to Router **MX-2**.

```
user@MX-1> monitor ethernet delay-measurement two-way mep 101 maintenance-domain md6
maintenance-association ma6 count 10
Two-way ETH-DM request to 00:90:69:0a:48:57, Interface ge-5/2/9.0
DMR received from 00:90:69:0a:48:57 Delay: 100 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 8 usec
```

```
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 111 usec Delay variation: 19 usec
DMR received from 00:90:69:0a:48:57 Delay: 110 usec Delay variation: 1 usec
DMR received from 00:90:69:0a:48:57 Delay: 119 usec Delay variation: 9 usec
DMR received from 00:90:69:0a:48:57 Delay: 122 usec Delay variation: 3 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 30 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

The counters are displayed as part of the MEP database on Router **MX-1** maintenance domain **MD6**.

```
user@MX-1> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 201, Direction: down, MAC address: 00:90:69:0a:43:94
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 894
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 0
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 10
    DMRs sent                                     : 0
    Valid DMRs received                           : 10
    Invalid DMRs received                         : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      101     00:90:69:0a:48:57       ok    ge-5/2/9.0
```

The collected MEP statistics are saved (up to 100 per remote MEP or per CFM session) and displayed as part of the MEP statistics on Router **MX-1**.

```
user@MX-1> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md6
```



```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
  CCMs sent                                     : 3154
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           : 0
  1DMs sent                                     : 0
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
  DMMs sent                                     : 10
  DMRs sent                                     : 0
  Valid DMRs received                           : 10
  Invalid DMRs received                         : 0

  Remote MEP identifier: 101
  Remote MAC address: 00:90:69:0a:48:57
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1                       100
      2                       92
      3                       92
      4                       111
      5                       110
      6                       119
      7                       122
      8                       92
      9                       92
      10                      108
    Average two-way delay          : 103 usec
    Average two-way delay variation: 8 usec
    Best case two-way delay        : 92 usec
    Worst case two-way delay       : 122 usec
```

The collected delay statistics are also saved (up to 100 per session) and displayed as part of the MEP delay statistics on Router **MX-1**.

```
user@MX-1> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1

  Remote MAC address: 00:90:69:0a:48:57
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
                         (usec)         (usec)
      1                       100
      2                       92
      3                       92
      4                       111
      5                       110
      6                       119
```

```
  7                       122
  8                       92
  9                       92
  10                      108
Average two-way delay          : 103 usec
Average two-way delay variation: 8 usec
Best case two-way delay        : 92 usec
Worst case two-way delay       : 122 usec
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 723
- Configuring MEP Interfaces to Support ETH-DM on page 618
- Triggering an ETH-DM Session on page 852
- Viewing ETH-DM Statistics on page 853
- Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 817
- Configuring ETH-DM with Untagged Interfaces on page 826

Example: Configuring Ethernet Frame Delay Measurements with Untagged Interfaces

Ethernet frame delay measurements are supported on untagged interfaces. All commands are the same as for tagged interfaces. Only the configurations are different. This section shows the untagged interface configurations for Routers **MX-1** and **MX-2**.

NOTE: **These are not complete router configurations.**

Untagged interface configuration for Router **MX-1**.

```
[edit]
interfaces {
  ge-5/0/0 {
    unit 0;
  }
  ge-5/2/9 {
    unit 0;
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        linktrace {
          path-database-size 255;
          age 10s;
        }
```

```
                maintenance-domain md6 {
                    level 6;
                    maintenance-association ma6 {
                        continuity-check {
                            interval 100ms;
                            hold-interval 1;
                        }
                        mep 201 {
                            interface ge-5/0/0;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
            }
        }
    }
}
```

Untagged interface configuration for Router **MX-2**.

```
[edit]
interfaces {
    ge-0/2/2 {
        unit 0;
    }
    ge-0/2/5 {
        unit 0;
    }
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                traceoptions {
                    file eoam_cfm.log size 1g files 2 world-readable;
                    flag all;
                }
                linktrace {
                    path-database-size 255;
                    age 10s;
                }
                maintenance-domain md6 {
                    level 6;
                    maintenance-association ma6 {
                        continuity-check {
                            interval 100ms;
                            hold-interval 1;
                        }
                        mep 101 {
                            interface ge-0/2/2;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
```

```
      }
    }
  }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 723
- Configuring MEP Interfaces to Support ETH-DM on page 618
- Triggering an ETH-DM Session on page 852
- Viewing ETH-DM Statistics on page 853
- Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 817
- Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 822

Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs

This example illustrates how to configure Ethernet frame loss measurement (ETH-LM) for single-tagged Loss Measurement Message (LMM)/Loss Measurement Reply (LMR) protocol data units (PDUs). By configuring ETH-LM, you can measure the Ethernet frame loss that occur in your network.

Requirements

This example uses the following hardware and software components:

- Two MX Series 3D Universal Edge Routers with Rev-B Dense Port Concentrators (DPCs)
- Junos OS Release 14.2 or later

Overview and Topology

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on Rev-B Dense Port Concentrators (DPCs) in MX Series routers. Additionally, the Y.1731 functionality supports ETH-LM only for an end-to-end connection that uses Virtual Private Wire Service (VPWS). This example illustrates how to configure ETH-LM for single-tagged LMM/LMR PDUs with input and output VLAN map configured as **swap**.

Figure 53 on page 829 shows the topology used in this example. VPWS service is configured between two MX Sereies routers, MX-PE1 and MX PE2.

Figure 53: VPWS Service Configured Between Two MX Series Routers

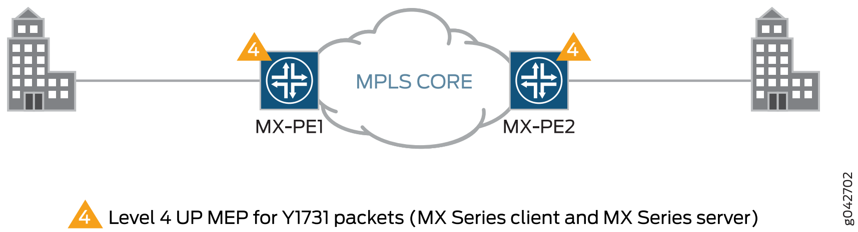

MX-PE1 router has two Ethernet interfaces, **ge-5/0/4** and **ge-5/1/9**. Virtual LAN (VLAN) is configured on **ge-5/0/4** and MPLS is configured on the **ge-5/1/9** interface. The **ge-5/0/4.11** interface is used to configure the Layer 2 virtual circuit with MX-PE2 router. The UP MEP, **mep 2**, is attached to the **ge-5/0/4.11** interface. The three-color policer firewall filter is also configured for the MX-PE1 router.

Similarly, MX-PE2 router has two Ethernet interfaces, **ge-8/0/8** and **ge-8/0/9**. Virtual LAN (VLAN) is configured on **ge-8/0/8** and MPLS is configured on the **ge-8/0/9** interface. The **ge-8/0/8.11** interface is used to configure the Layer 2 virtual circuit with MX-PE1 router. The UP MEP, **mep 1**, is attached to the **ge-8/0/8.11** interface. The three-color policer firewall filter is also configured for the MX-PE2 router.

Configuration

- Configuring Router PE1 on page 831
- Configuring Router PE2 on page 834

CLI Quick Configuration

To quickly configure ETH-LM for single-tagged LMM/LMR PDUs, copy the following commands, remove any line breaks, and then paste the commands into the CLI of each device.

On Router PE1:

```
[edit]
set interfaces ge-5/0/4 encapsulation flexible-ethernet-services
set interfaces ge-5/0/4 unit 11 encapsulation vlan-ccc
set interfaces ge-5/0/4 unit 11 layer2-policer input-three-color abc
set interfaces ge-5/0/4 unit 11 family ccc
set interfaces ge-5/1/9 enable
set interfaces ge-5/1/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-5/1/9 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 4.4.4.4/32
set interfaces ge-5/0/4 flexible-vlan-tagging
set interfaces ge-5/0/4 unit 11 vlan-id 2000
set interfaces ge-5/0/4 unit 11 input-vlan-map swap
set interfaces ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-5/0/4 unit 11 output-vlan-map swap
set routing-options router-id 4.4.4.4
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
set protocols oam ethernet connectivity-fault-management performance-monitoring
  delegate-server-processing
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 interface ge-5/0/4.11
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 direction up
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 remote-mep 1
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k
```

On Router PE2:

```
[edit]
set interfaces ge-8/0/8 encapsulation flexible-ethernet-services
set interfaces ge-8/0/8 unit 11 encapsulation vlan-ccc
set interfaces ge-8/0/8 unit 11 layer2-policer input-three-color abc
set interfaces ge-8/0/8 unit 11 family ccc
set interfaces ge-8/0/9 enable
set interfaces ge-8/0/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-8/0/9 unit 0 family mpls
set interfaces ae0 unit 0 family inet
set interfaces lo0 unit 0 family inet address 3.3.3.3/32
set interfaces ge-8/0/8 flexible-vlan-tagging
set interfaces ge-8/0/8 unit 11 vlan-id 2000
set interfaces ge-8/0/8 unit 11 input-vlan-map swap
set interfaces ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-8/0/8 unit 11 output-vlan-map swap
set routing-options router-id 3.3.3.3
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 1 interface ge-8/0/8.11
```


set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association ma mep 1 direction up
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association ma mep 1 remote-mep 2
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k

Configuring Router PE1

Step-by-Step Procedure

To configure Router PE1:

1. Configure the interfaces.

```
[edit]
user@PE1# edit interfaces
[edit interfaces]
user@PE1# set ge-5/0/4 encapsulation flexible-ethernet-services
user@PE1# set ge-5/0/4 unit 11 encapsulation vlan-ccc
user@PE1# set ge-5/0/4 unit 11 layer2-policer input-three-color abc
user@PE1# set ge-5/0/4 unit 11 family ccc
user@PE1# set ge-5/1/9 enable
user@PE1# set ge-5/1/9 unit 0 family inet address 12.1.1.1/24
user@PE1# set ge-5/1/9 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 4.4.4.4/32
```

2. Configure the VLAN.

```
[edit interfaces]
user@PE1# set ge-5/0/4 flexible-vlan-tagging
user@PE1# set ge-5/0/4 unit 11 vlan-id 2000
user@PE1# set ge-5/0/4 unit 11 input-vlan-map swap
user@PE1# set ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
user@PE1# set ge-5/0/4 unit 11 output-vlan-map swap
```

3. Configure the router identifier to identify the routing device.

```
[edit]
user@PE1# edit routing-options
[edit routing-options]
user@PE1# set router-id 4.4.4.4
```

4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE1# edit protocols
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

```
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable
```

5. Configure the Layer 2 circuit.

```
[edit protocols]
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
```

6. Configure the MEP.

```
[edit protocols]
user@PE1# set oam ethernet connectivity-fault-management
  performance-monitoring delegate-server-processing
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md level 4
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma continuity-check interval 1s
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 2 interface ge-5/0/4.11
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 2 direction up
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 2 remote-mep 1
```

7. Configure the firewall.

```
[edit]
user@PE1# edit firewall
[edit firewall]
user@PE1# set three-color-policer abc logical-interface-policer
user@PE1# set three-color-policer abc two-rate color-blind
user@PE1# set three-color-policer abc two-rate committed-information-rate 10m
user@PE1# set three-color-policer abc two-rate committed-burst-size 1500
user@PE1# set three-color-policer abc two-rate peak-information-rate 20m
user@PE1# set three-color-policer abc two-rate peak-burst-size 15k
```

8. Commit the configuration.

```
[edit]
user@PE1# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
interfaces {
  ge-5/0/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
```



```
      unit 11 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        input-vlan-map {
          swap;
          vlan-id 4094;
        }
        output-vlan-map swap;
        layer2-policer {
          input-three-color abc;
        }
        family ccc;
      }
    }
    ge-5/1/9 {
      enable;
      unit 0 {
        family inet {
          address 12.1.1.1/24;
        }
        family mpls;
      }
    }
    lo0 {
      unit 0 {
        family inet {
          address 4.4.4.4/32;
        }
      }
    }
  }
```

```
user@PE1# show protocols
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  l2circuit {
    neighbor 3.3.3.3 {
```

```
            interface ge-5/0/4.11 {
                virtual-circuit-id 1003;
                no-control-word;
            }
        }
    }
    oam {
        ethernet {
            connectivity-fault-management {
                performance-monitoring {
                    delegate-server-processing;
                }
                maintenance-domain md {
                    level 4;
                    maintenance-association ma {
                        continuity-check {
                            interval 1s;
                        }
                        mep 2 {
                            interface ge-5/0/4.11;
                            direction up;
                            remote-mep 1;
                        }
                    }
                }
            }
        }
    }
}
```

```
user@PE1# show routing-options
routing-options {
    router-id 4.4.4.4;
}
```

```
user@PE1# show firewall
firewall {
    three-color-policer abc {
        logical-interface-policer;
        two-rate {
            color-blind;
            committed-information-rate 10m;
            committed-burst-size 1500;
            peak-information-rate 20m;
            peak-burst-size 15k;
        }
    }
}
```

Configuring Router PE2

Step-by-Step Procedure

To configure Router PE2:

1. Configure the interfaces.

```
[edit]
```



```
user@PE2# edit interfaces
[edit interfaces]
user@PE2# set ge-8/0/8 encapsulation flexible-ethernet-services
user@PE2# set ge-8/0/8 unit 11 encapsulation vlan-ccc
user@PE2# set ge-8/0/8 unit 11 layer2-policer input-three-color abc
user@PE2# set ge-8/0/8 unit 11 family ccc
user@PE2# set ge-8/0/9 enable
user@PE2# set ge-8/0/9 unit 0 family inet address 12.1.1.1/24
user@PE2# set ge-8/0/9 unit 0 family mpls
user@PE2# set ae0 unit 0 family inet
user@PE2# set lo0 unit 0 family inet address 3.3.3.3/32
```

2. Configure the VLAN.

```
[edit interfaces]
user@PE2# set ge-8/0/8 flexible-vlan-tagging
user@PE2# set ge-8/0/8 unit 11 vlan-id 2000
user@PE2# set ge-8/0/8 unit 11 input-vlan-map swap
user@PE2# set ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
user@PE2# set ge-8/0/8 unit 11 output-vlan-map swap
```

3. Configure the router identifier to identify the routing device.

```
[edit]
user@PE2# edit routing-options
[edit routing-options]
user@PE2# set router-id 3.3.3.3
```

4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE2# edit protocols
[edit protocols]
user@PE2# set mpls interface all
user@PE2# set mpls interface fxp0.0 disable
user@PE2# set ospf area 0.0.0.0 interface all
user@PE2# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE2# set ldp interface all
user@PE2# set ldp interface fxp0.0 disable
```

5. Configure the Layer 2 circuit.

```
[edit protocols]
user@PE2# set l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id
  1003
user@PE2# set l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
```

6. Configure the MEP.

```
[edit protocols]
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md level 4
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma continuity-check interval 1s
```

```
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 1 interface ge-8/0/8.11
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 1 direction up
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 1 remote-mep 2
```

7. Configure the firewall.

```
[edit]
user@PE2# edit firewall
[edit firewall]
user@PE2# set three-color-policer abc logical-interface-policer
user@PE2# set three-color-policer abc two-rate color-blind
user@PE2# set three-color-policer abc two-rate committed-information-rate 10m
user@PE2# set three-color-policer abc two-rate committed-burst-size 1500
user@PE2# set three-color-policer abc two-rate peak-information-rate 20m
user@PE2# set three-color-policer abc two-rate peak-burst-size 15k
```

8. Commit the configuration.

```
[edit]
user@PE2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
interfaces {
  ge-8/0/8 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-id 2000;
      input-vlan-map {
        swap;
        vlan-id 4094;
      }
      output-vlan-map swap;
      layer2-policer {
        input-three-color abc;
      }
      family ccc;
    }
  }
  ge-8/0/9 {
    unit 0 {
      family inet {
        address 12.1.1.2/24;
      }
```



```
            family mpls;
        }
    }
    ae0 {
        unit 0 {
            family inet;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 3.3.3.3/32;
            }
        }
    }
}
```

user@PE2# **show protocols**

```
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    l2circuit {
        neighbor 4.4.4.4 {
            interface ge-8/0/8.11 {
                virtual-circuit-id 1003;
                no-control-word;
            }
        }
    }
    oam {
        ethernet {
            connectivity-fault-management {
                maintenance-domain md {
                    level 4;
                    maintenance-association ma {
                        continuity-check {
                            interval 1s;
                        }
```



```
                mep 1 {
                    interface ge-8/0/8.11;
                    direction up;
                    remote-mep 2;
                }
            }
        }
      }
    }
  }
}
```

```
user@PE2# show routing-options
routing-options {
  router-id 3.3.3.3;
}
```

```
user@PE2# show firewall
firewall {
  three-color-policer abc {
    logical-interface-policer;
    two-rate {
      color-blind;
      committed-information-rate 10m;
      committed-burst-size 1500;
      peak-information-rate 20m;
      peak-burst-size 15k;
    }
  }
}
```

Verification

To start monitoring the Ethernet frame loss, issue the **monitor ethernet loss-measurement maintenance-domain md maintenance-association ma mep 1** command. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the **monitor ethernet loss-measurement** command. You can also issue the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command to display ETH-LM statistics.

Viewing ETH-LM

Purpose View the ETH-LM statistics.

Action From operational mode, enter the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command.

```
user@PE1> show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11
Interface name: ge-5/0/4.11 , Interface status: Active, Link status: Up
  Maintenance domain name: md, Format: string, Level: 4
```



```
Maintenance association name: ma, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
MEP identifier: 2, Direction: up, MAC address: 00:24:dc:9b:96:76
MEP status: running
Defects:
  Remote MEP not receiving CCM                  : no
  Erroneous CCM received                        : no
  Cross-connect CCM received                    : no
  RDI sent by some MEP                          : no
  Some remote MEP's MAC in error state          : no
Statistics:
  CCMs sent                                     : 36
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           : 0
  1DMs sent                                     : 0
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
  Out of sync 1DMs received                     : 0
  DMMs sent                                     : 0
  Valid DMMs received                           : 0
  Invalid DMMs received                         : 0
  DMRs sent                                     : 0
  Valid DMRs received                           : 0
  Invalid DMRs received                         : 0
  LMMs sent                                     : 10
  Valid LMMs received                           : 0
  Invalid LMMs received                         : 0
  LMRs sent                                     : 0
  Valid LMRs received                           : 10
  Invalid LMRs received                         : 0
  SLMs sent                                     : 0
  Valid SLMs received                           : 0
  Invalid SLMs received                         : 0
  SLRs sent                                     : 0
  Valid SLRs received                           : 0
  Invalid SLRs received                         : 0
Remote MEP count: 1
  Identifier    MAC address        State    Interface
      1     00:05:85:76:e5:30       ok   ge-5/0/4.11
```

Meaning The Ethernet interface details and statistics are displayed. This output indicates that the **ge-5/0/4.11** interface is active and its link status is **up**. Its maintenance domain name is **md** and its level is **4**. The MEP identifier of the **ge-5/0/4.11** interface is indicated as **2** and its direction is **up**. Under the statistics section, the output indicates that **10** LMMs were sent and **10** valid LMRs were received by the interface.

Related Documentation

- Ethernet Frame Loss Measurement Overview on page 729
- Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs on page 840

Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs

This example illustrates how to configure Ethernet frame loss measurement (ETH-LM) for dual-tagged Loss Measurement Message (LMM)/Loss Measurement Reply (LMR) protocol data units (PDUs). By configuring ETH-LM, you can measure the Ethernet frame loss that occur in your network.

Requirements

This example uses the following hardware and software components:

- Two MX Series 3D Universal Edge Routers with Rev-B Dense Port Concentrators (DPCs)
- Junos OS Release 14.2 or later

Overview and Topology

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on Rev-B Dense Port Concentrators (DPCs) in MX Series routers. Additionally, the Y.1731 functionality supports ETH-LM only for an end-to-end connection that uses Virtual Private Wire Service (VPWS). This example illustrates how to configure ETH-LM for dual tagged LMM/LMR PDUs with input and output VLAN map configured as **swap-swap**.

Figure 54 on page 840 shows the topology used in this example. VPWS service is configured between two MX Sereies routers, MX-PE1 and MX PE2.

Figure 54: VPWS Service Configured Between Two MX Series Routers

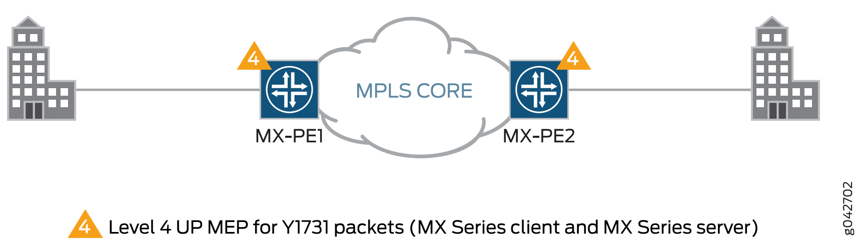

MX-PE1 router has two Ethernet interfaces, **ge-5/0/4** and **ge-5/1/9**. Virtual LAN (VLAN) is configured on **ge-5/0/4** and MPLS is configured on the **ge-5/1/9** interface. The **ge-5/0/4.11** interface is used to configure the Layer 2 virtual circuit with MX-PE2 router.

The UP MEP, **mep 2**, is attached to the **ge-5/0/4.11** interface. The three-color policer firewall filter is also configured for the MX-PE1 router.

Similarly, MX-PE2 router has two Ethernet interfaces, **ge-8/0/8** and **ge-8/0/9**. Virtual LAN (VLAN) is configured on **ge-8/0/8** and MPLS is configured on the **ge-8/0/9** interface. The **ge-8/0/8.11** interface is used to configure the Layer 2 virtual circuit with MX-PE1 router. The UP MEP, **mep 1**, is attached to the **ge-8/0/8.11** interface. The three-color policer firewall filter is also configured for the MX-PE2 router.

Configuration

- Configuring Router PE1 on page 843
- Configuring Router PE2 on page 846

CLI Quick Configuration

To quickly configure ETH-LM for dual tagged LMM/LMR PDUs, copy the following commands, remove any line breaks, and then paste the commands into the CLI of each device.

On Router PE1:

```
[edit]
set interfaces ge-5/0/4 encapsulation flexible-ethernet-services
set interfaces ge-5/0/4 unit 11 encapsulation vlan-ccc
set interfaces ge-5/0/4 unit 11 layer2-policer input-three-color abc
set interfaces ge-5/0/4 unit 11 family ccc
set interfaces ge-5/1/9 enable
set interfaces ge-5/1/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-5/1/9 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 4.4.4.4/32
set interfaces ge-5/0/4 flexible-vlan-tagging
set interfaces ge-5/0/4 unit 11 vlan-tags outer 2000 inner 1000
set interfaces ge-5/0/4 unit 11 input-vlan-map swap-swap
set interfaces ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-5/0/4 unit 11 input-vlan-map inner-vlan-id 4093
set interfaces ge-5/0/4 unit 11 output-vlan-map swap-swap
set routing-options router-id 4.4.4.4
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
set protocols oam ethernet connectivity-fault-management performance-monitoring
  delegate-server-processing
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 interface ge-5/0/4.11
set protocols oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 direction up
```

```
set protocols oam ethernet connectivity-fault-management maintenance-domain md
   maintenance-association ma mep 2 remote-mep 1
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k
```

On Router PE2:

```
[edit]
set interfaces ge-8/0/8 encapsulation flexible-ethernet-services
set interfaces ge-8/0/8 unit 11 encapsulation vlan-ccc
set interfaces ge-8/0/8 unit 11 layer2-policer input-three-color abc
set interfaces ge-8/0/8 unit 11 family ccc
set interfaces ge-8/0/9 enable
set interfaces ge-8/0/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-8/0/9 unit 0 family mpls
set interfaces ae0 unit 0 family inet
set interfaces lo0 unit 0 family inet address 3.3.3.3/32
set interfaces ge-8/0/8 flexible-vlan-tagging
set interfaces ge-8/0/8 unit 11 vlan-tags outer 2000 inner 1000
set interfaces ge-8/0/8 unit 11 input-vlan-map swap-swap
set interfaces ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-8/0/8 unit 11 input-vlan-map inner-vlan-id 4093
set interfaces ge-8/0/8 unit 11 output-vlan-map swap-swap
set routing-options router-id 3.3.3.3
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
set protocols oam ethernet connectivity-fault-management maintenance-domain md
   level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md
   maintenance-association ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md
   maintenance-association ma mep 1 interface ge-8/0/8.11
set protocols oam ethernet connectivity-fault-management maintenance-domain md
   maintenance-association ma mep 1 direction up
set protocols oam ethernet connectivity-fault-management maintenance-domain md
   maintenance-association ma mep 1 remote-mep 2
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k
```


Configuring Router PE1

Step-by-Step Procedure

To configure Router PE1:

1. Configure the interfaces.

```
[edit]
user@PE1# edit interfaces
[edit interfaces]
user@PE1# set ge-5/0/4 encapsulation flexible-ethernet-services
user@PE1# set ge-5/0/4 unit 11 encapsulation vlan-ccc
user@PE1# set ge-5/0/4 unit 11 layer2-policer input-three-color abc
user@PE1# set ge-5/0/4 unit 11 family ccc
user@PE1# set ge-5/1/9 enable
user@PE1# set ge-5/1/9 unit 0 family inet address 12.1.1.1/24
user@PE1# set ge-5/1/9 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 4.4.4.4/32
```

2. Configure the VLAN.

```
[edit interfaces]
user@PE1# set ge-5/0/4 flexible-vlan-tagging
user@PE1# set ge-5/0/4 unit 11 vlan-tags outer 2000 inner 1000
user@PE1# set ge-5/0/4 unit 11 input-vlan-map swap-swap
user@PE1# set ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
user@PE1# set ge-5/0/4 unit 11 input-vlan-map inner-vlan-id 4093
user@PE1# set ge-5/0/4 unit 11 output-vlan-map swap-swap
```

3. Configure the router identifier to identify the routing device.

```
[edit]
user@PE1# edit routing-options
[edit routing-options]
user@PE1# set router-id 4.4.4.4
```

4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE1# edit protocols
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable
```

5. Configure the Layer 2 circuit.

```
[edit protocols]
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
```

6. Configure the MEP.

```
[edit protocols]
user@PE1# set oam ethernet connectivity-fault-management
  performance-monitoring delegate-server-processing
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md level 4
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma continuity-check interval 1s
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 2 interface ge-5/0/4.11
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 2 direction up
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 2 remote-mep 1
```

7. Configure the firewall.

```
[edit]
user@PE1# edit firewall
[edit firewall]
user@PE1# set three-color-policer abc logical-interface-policer
user@PE1# set three-color-policer abc two-rate color-blind
user@PE1# set three-color-policer abc two-rate committed-information-rate 10m
user@PE1# set three-color-policer abc two-rate committed-burst-size 1500
user@PE1# set three-color-policer abc two-rate peak-information-rate 20m
user@PE1# set three-color-policer abc two-rate peak-burst-size 15k
```

8. Commit the configuration.

```
[edit]
user@PE1# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
interfaces {
  ge-5/0/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-tags outer 2000 inner 1000;
      input-vlan-map {
        swap-swap;
        vlan-id 4094;
        inner-vlan-id 4093;
      }
      output-vlan-map swap-swap;
      layer2-policer {
```

```
            input-three-color abc;
          }
          family ccc;
        }
      }
      ge-5/1/9 {
        enable;
        unit 0 {
          family inet {
            address 12.1.1.1/24;
          }
          family mpls;
        }
      }
      lo0 {
        unit 0 {
          family inet {
            address 4.4.4.4/32;
          }
        }
      }
    }
```

```
user@PE1# show protocols
protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  l2circuit {
    neighbor 3.3.3.3 {
      interface ge-5/0/4.11 {
        virtual-circuit-id 1003;
        no-control-word;
      }
    }
  }
  oam {
    ethernet {
      connectivity-fault-management {
```

```
            performance-monitoring {
              delegate-server-processing;
            }
            maintenance-domain md {
              level 4;
              maintenance-association ma {
                continuity-check {
                  interval 1s;
                }
                mep 2 {
                  interface ge-5/0/4.11;
                  direction up;
                  remote-mep 1;
                }
              }
            }
          }
        }
      }
    }
```

```
user@PE1# show routing-options
routing-options {
  router-id 4.4.4.4;
}
```

```
user@PE1# show firewall
firewall {
  three-color-policer abc {
    logical-interface-policer;
    two-rate {
      color-blind;
      committed-information-rate 10m;
      committed-burst-size 1500;
      peak-information-rate 20m;
      peak-burst-size 15k;
    }
  }
}
```

Configuring Router PE2

Step-by-Step Procedure

To configure Router PE2:

1. Configure the interfaces.

```
[edit]
user@PE2# edit interfaces
[edit interfaces]
user@PE2# set ge-8/0/8 encapsulation flexible-ethernet-services
user@PE2# set ge-8/0/8 unit 11 encapsulation vlan-ccc
user@PE2# set ge-8/0/8 unit 11 layer2-policer input-three-color abc
user@PE2# set ge-8/0/8 unit 11 family ccc
user@PE2# set ge-8/0/9 enable
user@PE2# set ge-8/0/9 unit 0 family inet address 12.1.1.1/24
user@PE2# set ge-8/0/9 unit 0 family mpls
```



```
user@PE2# set ae0 unit 0 family inet
user@PE2# set lo0 unit 0 family inet address 3.3.3.3/32
```

2. Configure the VLAN.

```
[edit interfaces]
user@PE2# set ge-8/0/8 flexible-vlan-tagging
user@PE2# set ge-8/0/8 unit 11 vlan-tags outer 2000 inner 1000
user@PE2# set ge-8/0/8 unit 11 input-vlan-map swap-swap
user@PE2# set ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
user@PE2# set ge-8/0/8 unit 11 input-vlan-map inner-vlan-id 4093
user@PE2# set ge-8/0/8 unit 11 output-vlan-map swap-swap
```

3. Configure the router identifier to identify the routing device.

```
[edit]
user@PE2# edit routing-options
[edit routing-options]
user@PE2# set router-id 3.3.3.3
```

4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE2# edit protocols
[edit protocols]
user@PE2# set mpls interface all
user@PE2# set mpls interface fxp0.0 disable
user@PE2# set ospf area 0.0.0.0 interface all
user@PE2# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE2# set ldp interface all
user@PE2# set ldp interface fxp0.0 disable
```

5. Configure the Layer 2 circuit.

```
[edit protocols]
user@PE2# set l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id
  1003
user@PE2# set l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
```

6. Configure the MEP.

```
[edit protocols]
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md level 4
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma continuity-check interval 1s
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 1 interface ge-8/0/8.11
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 1 direction up
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain
  md maintenance-association ma mep 1 remote-mep 2
```

7. Configure the firewall.

```
[edit]
user@PE2# edit firewall
[edit firewall]
user@PE2# set three-color-policer abc logical-interface-policer
user@PE2# set three-color-policer abc two-rate color-blind
user@PE2# set three-color-policer abc two-rate committed-information-rate 10m
user@PE2# set three-color-policer abc two-rate committed-burst-size 1500
user@PE2# set three-color-policer abc two-rate peak-information-rate 20m
user@PE2# set three-color-policer abc two-rate peak-burst-size 15k
```

8. Commit the configuration.

```
[edit]
user@PE2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-options**, and **show firewall** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
interfaces {
  ge-8/0/8 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-tags outer 2000 inner 1000;
      input-vlan-map {
        swap-swap;
        vlan-id 4094;
        inner-vlan-id 4093;
      }
      output-vlan-map swap-swap;
      layer2-policer {
        input-three-color abc;
      }
      family ccc;
    }
  }
  ge-8/0/9 {
    unit 0 {
      family inet {
        address 12.1.1.2/24;
      }
      family mpls;
    }
  }
  ae0 {
    unit 0 {
      family inet;
```

```
      }
    }
    lo0 {
      unit 0 {
        family inet {
          address 3.3.3.3/32;
        }
      }
    }
  }

  user@PE2# show protocols
  protocols {
    mpls {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
    ospf {
      area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
          disable;
        }
      }
    }
    ldp {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
    l2circuit {
      neighbor 4.4.4.4 {
        interface ge-8/0/8.11 {
          virtual-circuit-id 1003;
          no-control-word;
        }
      }
    }
    oam {
      ethernet {
        connectivity-fault-management {
          maintenance-domain md {
            level 4;
            maintenance-association ma {
              continuity-check {
                interval 1s;
              }
              mep 1 {
                interface ge-8/0/8.11;
                direction up;
                remote-mep 2;
              }
            }
```

```
        }
      }
    }
  }
}
```

user@PE2# **show routing-options**

```
routing-options {
  router-id 3.3.3.3;
}
```

user@PE2# **show firewall**

```
firewall {
  three-color-policer abc {
    logical-interface-policer;
    two-rate {
      color-blind;
      committed-information-rate 10m;
      committed-burst-size 1500;
      peak-information-rate 20m;
      peak-burst-size 15k;
    }
  }
}
```

Verification

To start the Ethernet frame loss measurement session, issue the **monitor ethernet loss-measurement maintenance-domain md maintenance-association ma mep 1** command. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the **monitor ethernet loss-measurement** command. You can also issue the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command to display ETH-LM statistics.

Viewing ETH-LM

Purpose View the ETH-LM statistics.

Action From operational mode, enter the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command.

```
user@PE1> show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11
Interface name: ge-5/0/4.11 , Interface status: Active, Link status: Up
  Maintenance domain name: md, Format: string, Level: 4
    Maintenance association name: ma, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: none, Port status TLV: none
    Connection Protection TLV: no
    MEP identifier: 2, Direction: up, MAC address: 00:24:dc:9b:96:76
    MEP status: running
    Defects:
```



```
    Remote MEP not receiving CCM                 : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                         : no
    Some remote MEP's MAC in error state         : no
  Statistics:
    CCMs sent                                    : 59
    CCMs received out of sequence                : 0
    LBMs sent                                    : 0
    Valid in-order LBRs received                 : 0
    Valid out-of-order LBRs received             : 0
    LBRs received with corrupted data            : 0
    LBRs sent                                    : 0
    LTMs sent                                    : 0
    LTMs received                                : 0
    LTRs sent                                    : 0
    LTRs received                                : 0
    Sequence number of next LTM request          : 0
    1DMs sent                                    : 0
    Valid 1DMs received                          : 0
    Invalid 1DMs received                        : 0
    Out of sync 1DMs received                    : 0
    DMMs sent                                    : 0
    Valid DMMs received                          : 0
    Invalid DMMs received                        : 0
    DMRs sent                                    : 0
    Valid DMRs received                          : 0
    Invalid DMRs received                        : 0
    LMMs sent                                    : 10
    Valid LMMs received                          : 0
    Invalid LMMs received                        : 0
    LMRs sent                                    : 0
    Valid LMRs received                          : 10
    Invalid LMRs received                        : 0
    SLMs sent                                    : 0
    Valid SLMs received                          : 0
    Invalid SLMs received                        : 0
    SLRs sent                                    : 0
    Valid SLRs received                          : 0
    Invalid SLRs received                        : 0
  Remote MEP count: 1
    Identifier     MAC address        State    Interface
        1      00:05:85:76:e5:30       ok   ge-5/0/4.11
```

Meaning

The Ethernet interface details and statistics are displayed. This output indicates that the **ge-5/0/4.11** interface is active and its link status is **up**. Its maintenance domain name is **md** and its level is **4**. The MEP identifier of the **ge-5/0/4.11** interface is indicated as **2** and its direction is **up**. Under the statistics section, the output indicates that **10** LMMs were sent and **10** valid LMRs were received by the interface.

Related Documentation

- Ethernet Frame Loss Measurement Overview on page 729
- Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs on page 828

Triggering an Ethernet Frame Delay Measurements Session

Before Ethernet frame delay measurement statistics can be displayed, they must be collected. To trigger Ethernet frame delay measurement, use the **monitor ethernet delay-measurement (one-way | two-way) (*remote-mac-address* | mep *identifier*) maintenance-domain *name* maintenance-association *ma-id* [count *count*] [wait *time*]** operational command.

The fields for this command are described in Table 81 on page 852.

Table 81: Monitor Ethernet Delay Command Parameters

Parameter	Parameter Range	Description
one-way or **two-way**	NA	Perform a one-way or two-way (round-trip) delay measurement.
remote-mac-address	Unicast MAC address	Send delay measurement frames to the destination unicast MAC address (use the format xx:xx:xx:xx:xx:xx). Multicast MAC addresses are not supported.
mep *identifier*	1–8191	The MEP identifier to use for the measurement. The discovered MAC address for this MEP identifier is used.
maintenance-domain *name*	Existing MD name	Specifies an existing maintenance domain (MD) to use for the measurement.
maintenance-association *ma-id*	Existing MA identifier	Specifies an existing maintenance association (MA) identifier to use for the measurement.
count *count*	1–65535 (default: 10)	(Optional) Specifies the number of Ethernet frame delay frames to send. The default is 10.
wait *time*	1–255 seconds (default: 1)	(Optional) Specifies the number of seconds to wait between frames. The default is 1 second.

If you attempt to monitor delays to a nonexistent MAC address, you must exit the application manually using **^C**:

```
user@host> monitor ethernet delay-measurement two-way 00:11:22:33:44:55
Two-way ETH-DM request to 00:11:22:33:44:55, Interface ge-5/2/9.0
^C
--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 0
Average delay: 0 usec, Average delay variation: 0 usec
Best case delay: 0 usec, Worst case delay: 0 usec
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 723
- Configuring MEP Interfaces to Support ETH-DM on page 618
- Viewing ETH-DM Statistics on page 853

- Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 817
- Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 822
- Configuring ETH-DM with Untagged Interfaces on page 826

Viewing Ethernet Frame Delay Measurements Statistics

Once Ethernet frame delay measurement statistics have been collected, they can be displayed.

To retrieve the last 100 Ethernet frame delay measurement statistics per remote MEP or per CFM session, two types of **show** commands are provided:

- For all OAM frame counters and Ethernet frame delay measurement statistics
- For Ethernet frame delay measurement statistics only

To retrieve all Ethernet frame delay measurement statistics for a given session, use the **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *name* maintenance-association *name* [local-mep *identifier*] [remote-mep *identifier*] [count *count*]** command.

To retrieve only Ethernet frame delay measurement statistics for a given session, use the **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain *name* maintenance-association *name* [local-mep *identifier*] [remote-mep *identifier*] [count *count*]** command.

NOTE: The only difference in the two commands is the use of the **mep-statistics** and **delay-statistics** keyword.

The fields for these commands are described in Table 82 on page 853.

Table 82: Show Ethernet Delay Command Parameters

Parameter	Parameter Range	Description
maintenance-domain *name*	Existing MD name	Specifies an existing maintenance domain (MD) to use.
maintenance-association *ma-id*	Existing MA identifier	Specifies an existing maintenance association (MA) identifier to use.
local-mep *identifier*	1–8191	When a MEP has been specified, display statistics only for the local MEP.
remote-mep *identifier*	1–8191	When a MEP has been specified, display statistics only for the discovered MEP.
count *count*	1–100 (default:100)	The number of entries to display in the results table. By default, all 100 entries are displayed if they exist.

NOTE: For each MEP, you will see frame counters for sent and received Ethernet frame delay measurement frames whenever MEP statistics are displayed.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 723
- Configuring MEP Interfaces to Support ETH-DM on page 618
- Triggering an ETH-DM Session on page 852
- Configuring One-Way ETH-DM with Single-Tagged Interfaces on page 817
- Configuring Two-Way ETH-DM with Single-Tagged Interfaces on page 822
- Configuring ETH-DM with Untagged Interfaces on page 826

CHAPTER 34

Configuring Ethernet Ring Protection

Ethernet Ring Protection

Link failure is often an unavoidable part of networking. However, there are methods of improving the reliability of a router or bridge network even when link failures occur. For example, SONET/SDH seal-healing rings are frequently used to add a level of robustness to router networks. This ring protection switching is now extended to Ethernet links. You can configure Ethernet ring protection for a series of two or more systems so that if one link fails, traffic is rerouted around the failure on the ring.

The basic idea of Ethernet ring protection is to use one specific link to protect the whole ring. This special link is the ring protection link (RPL). When all links are up and running, the RPL blocks traffic and remains idle. The RPL itself is controlled by the designated RPL owner node. There is only one RPL owner node on the ring and the RPL owner node is responsible for blocking the RPL interface under normal operating conditions. However, if a link failure occurs on the ring, the RPL owner node is responsible for unblocking the RPL interface and protection–switching the traffic on the alternate path around the ring. An Ethernet ring automatic protection switching (R-APS) messaging protocol coordinates the protection activities of all nodes on the ring. The APS blocks traffic over the failed link and unblocks traffic over the RPL.

When the failed link is repaired, the traffic reverts to its normal pattern. That is, the RPL owner blocks the RPL link and unblocks traffic over the cleared link.

Two or more nodes form a ring. Links between the nodes form a chain, with the last node also connecting the first. Every ring node therefore has two ports related to the ring, one in each direction. In this chapter, these directions are referred to as east and west.

Every node on the ring is one of two types:

- RPL owner node—This node owns the RPL and blocks or unblocks the RPL as conditions require. This node initiates the R-APS message.
- Normal node—All other nodes on the ring (that is, those that are not the RPL owner node) operate as normal nodes and have no special role on the ring.

In addition to roles, each node on the Ethernet ring can be in one of several states:

- Init—The node is not yet participating in the ring.
- Idle—The node is performing normally (there is no link failure on the ring). In this state, traffic is unblocked on both ring ports, except for the RPL owner node, which blocks the RPL port (the other RPL owner port is unblocked).
- Protection—When a failure occurs on the ring, a normal node will have traffic blocked on the ring port that connects to the failed link. The RPL owner, if it is not at one end of the failed link, will then unblock the RPL port so both ports are active.

NOTE: **The R-APS protocol does not detect the number of RPL owner nodes configured on the ring. You must configure only one RPL and RPL owner per ring or protection switching will not work properly.**

Ethernet ring protection only works when one link on the ring fails. Multiple link failures will break the ring and cause protection switching to fail.

Several restrictions apply to Ethernet ring protection:

- The Ethernet ring protection configured as a single instance only works at the physical level (adjacent nodes must be directly connected). The ring protection operates at the interface (port) level and not at the VLAN level.
- Manual (command-based) switching to protection mode is not supported.
- Nonrevertive switching is not supported. When the link failure is cleared, traffic always returns to normal operation.
- The interconnection of multiple rings for protection purposes is not supported.

You can configure Ethernet ring protection to optomize traffic load-balancing by using multiple ring instances. For more information about multiple ring instances, see “Ethernet Ring Protection Using Ring Instances for Load Balancing” on page 857

Related Documentation

- *Ethernet OAM Feature Guide for MX Series Routers*
- Example: Configuring Ethernet Ring Protection for MX Series Routers on page 858
- Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 882
- Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 884
- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857

- Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers on page 864

Ethernet Ring Protection Using Ring Instances for Load Balancing

Juniper Network MX Series 3D Universal Edge Routers support Ethernet ring protection (ERP) to help achieve high reliability and network stability. ERP is used in router or bridge networks to protect against link failure. A single-ring topology is configured that uses one specific link called a ring protection link (RPL) to protect the whole ring. When all links are up and running, the RPL blocks traffic and remains idle. However, if a link fails, the RPL routes traffic to bypass the failure on the ring.

NOTE: **To learn how ERP works in a single-ring topology, see “Ethernet Ring Protection” on page 855.**

MX Series routers now support ERP ring instances. Whereas traffic in a single-ring topology follows the same path, traffic within ring instances allows some traffic to pass through one path while other traffic can follow a different path. Dividing traffic in this way supports traffic load balancing in the physical ring.

Ring instances are like traffic channels that contain different sets of virtual LANS (VLANs). A ring instance is responsible for the protection of a subset of VLANs that transport traffic over the physical ring. When ring instances are configured for the ring, each ring instance should have its own RPL owner, an east and a west interface, and a ring protection link end.

Each ring instance has a control channel and a specific data channel. A data channel is a group of bridge domain VLAN IDs. All VLAN IDs within the same ring interface share the same data-forwarding properties controlled by the ERP. If no data channel is defined in the ring configuration, ERP will only operate on the physical link instead of as a ring instance using logical links.

When operating ERP in a topology with other protocols, the following considerations should be observed:

- If a physical interface is part of an Ethernet ring, it cannot be configured for Spanning Tree Protocol (STP) or Multiple Spanning Tree Protocol (MSTP).
- ERP and Per-VLAN Spanning Tree (PVST) can be configured on the same topology as long as PVST doesn't share the same VLAN with any Ethernet ring instance configured on the physical port.
- If ERP is configured only as a physical ring instance (a ring without a data channel) in a topology also configured for PVST, ERP checks the PVST configuration on two ring interfaces and automatically creates a data channel excluding VLANs used by PVST.

Related Documentation

- Ethernet Ring Protection on page 855

Example: Configuring Ethernet Ring Protection for MX Series Routers

This example configures Ethernet ring protection for three MX Series router nodes:

Example Topology

The links connecting the three MX Series routers are shown in Figure 55 on page 858.

Figure 55: Ethernet Ring Protection Example Nodes

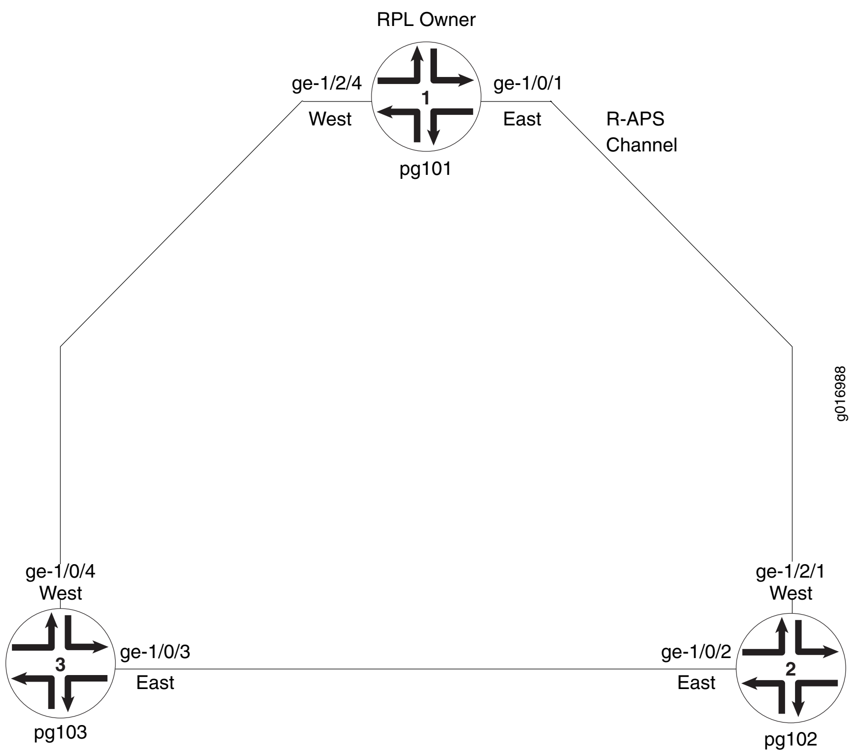

This example uses the following topology details for Ethernet ring protection:

- Router 1 is the RPL owner. The node identification for Router 1 is MAC address **00:01:01:00:00:01**.
- The RPL link is **ge-1/0/1.1** (this is also the R-APS messaging control channel).
- Traffic flows among the nodes in the configured bridge domains. (That is, only the control channels are configured.)

- Router 1's east control channel interface is **ge-1/0/1.1** (the RPL) and the west control channel interface is **ge-1/2/4.1**. The protection group is **pg101**.
- Router 2's east control channel interface is **ge-1/0/2.1** (the RPL) and the west control channel interface is **ge-1/2/1.1**. The protection group is **pg102**.
- Router 3's east control channel interface is **ge-1/0/3.1** (the RPL) and the west control channel interface is **ge-1/0/4.1**. The protection group is **pg103**.

NOTE: Although not strictly required for physical ring protection, this example configures Ethernet OAM with MEPs.

Router 1 (RPL Owner) Configuration

To configure Router 1 (the RPL owner):

1. Configure the interfaces:

```
[edit]
interfaces {
  ge-1/0/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  ge-1/2/4 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  irb {
    unit 0 {
      family inet {
        address address 192.1.1.11/24;
      }
    }
  }
}
```

2. Configure the bridge domain:

```
[edit]
bridge-domains {
  bd1 {
    domain-type bridge;
    vlan-id 100;
    interface ge-1/2/4.1;
    interface ge-1/0/1.1;
```

```
        routing-interface irb.0;
      }
    }
```

3. Configure the Ethernet ring protection group:

```
[edit]
protocols {
  protection-group {
    ethernet-ring pg101 {
      node-id 00:01:01:00:00:01;
      ring-protection-link-owner;
      east-interface {
        control-channel ge-1/0/1.1;
        }
      ring-protection-link-end;
      west-interface {
        control-channel ge-1/2/4.1;
      }
    }
  }
}
```

4. Configure Ethernet OAM:

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile rmep-defaults {
          default-action {
            interface-down;
          }
        }
        maintenance-domain d1 {
          level 0;
          maintenance-association 100 {
            mep 1 {
              interface ge-1/0/1;
              remote-mep 2 {
                action-profile rmep-defaults;
              }
            }
          }
        }
        maintenance-domain d2 {
          level 0;
          maintenance-association 100 {
            mep 1 {
              interface ge-1/2/4;
              remote-mep 2 {
                action-profile rmep-defaults;
              }
            }
          }
        }
```



```
            }
          }
        }
      }
```

Router 2 Configuration

To configure Router 2:

1. Configure the interfaces:

```
[edit]
interfaces {
  ge-1/0/2 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  ge-1/2/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  irb {
    unit 0 {
      family inet {
        address address 192.1.1.22/24;
      }
    }
  }
}
```

2. Configure the bridge domain:

```
[edit]
bridge-domains {
  bd1 {
    domain-type bridge;
    vlan-id 100;
    interface ge-1/2/1.1;
    interface ge-1/0/2.1;
    routing-interface irb.0;
  }
}
```

3. Configure the Ethernet protection group:

```
[edit]
protocols {
  protection-group {
    ethernet-ring pg102 {
```

```
            node-id 00:22:22:22:22:22;
            east-interface {
                control-channel ge-1/0/2.1;
            }
            west-interface {
                control-channel ge-1/2/1.1;
            }
        }
    }
}
```

4. Configure Ethernet OAM:

```
[edit]
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
                maintenance-domain d1 {
                    level 0;
                    maintenance-association 100 {
                        mep 2 {
                            interface ge-1/2/1;
                            remote-mep 1 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
                maintenance-domain d3 {
                    level 0;
                    maintenance-association 100 {
                        mep 1 {
                            interface ge-1/0/2;
                            remote-mep 2 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Router 3 Configuration

To configure Router 3:

1. Configure the interfaces:


```
[edit]
interfaces {
   ge-1/0/4 {
      vlan-tagging;
      encapsulation flexible-ethernet-services;
      unit 1 {
         encapsulation vlan-bridge;
         vlan-id 100;
      }
   }
   ge-1/0/3 {
      vlan-tagging;
      encapsulation flexible-ethernet-services;
      unit 1 {
         encapsulation vlan-bridge;
         vlan-id 100;
      }
   }
   irb {
      unit 0 {
         family inet {
            address 192.1.1.33/24;
         }
      }
   }
}
```

2. Configure the bridge domain:

```
[edit]
bridge-domains {
   bd1 {
      domain-type bridge;
      vlan-id 100;
      interface ge-1/0/4.1;
      interface ge-1/0/3.1;
      routing-interface irb.0;
   }
}
```

3. Configure the Ethernet protection group:

```
[edit]
protocols {
   protection-group {
      ethernet-ring pg103 {
         node-id 00:33:33:33:33:33;
         east-interface {
            control-channel ge-1/0/3.1;
         }
         west-interface {
            control-channel ge-1/0/4.1;
         }
      }
   }
}
```

4. Configure Ethernet OAM:

```
[edit]
protocols {
   oam {
      ethernet {
         connectivity-fault-management {
            action-profile rmep-defaults {
               default-action {
                  interface-down;
               }
            }
            maintenance-domain d2 {
               level 0;
               maintenance-association 100 {
                  mep 2 {
                     interface ge-1/0/4;
                     remote-mep 1 {
                        action-profile rmep-defaults;
                     }
                  }
               }
            }
            maintenance-domain d3 {
               level 0;
               maintenance-association 100 {
                  mep 2 {
                     interface ge-1/0/3;
                     remote-mep 1 {
                        action-profile rmep-defaults;
                     }
                  }
               }
            }
         }
      }
   }
}
```

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Ring Protection on page 855
- Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 882
- Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 884

Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers

MX Series routers support Ethernet ring protection (ERP) to help achieve high reliability and network stability. ERP is used in router or bridge networks to protect against link failure. A single-ring topology is configured that uses one specific link called a ring protection link (RPL) to protect the whole ring. When all links are up and running, the

RPL blocks traffic and remains idle. However, if a link fails, the RPL routes traffic to bypass the failure on the ring.

MX Series routers now support ERP ring instances. Whereas traffic in a ring topology follows the same path, traffic within a ring instance uses data channels to allow some traffic to pass through one path while other traffic can follow a different one. Dividing traffic in this way supports traffic load-balancing in the ring.

This example describes how to use ERP with ring instances to load-balance traffic while still providing network protection from link failure:

Requirements

This example uses the following hardware and software components:

- Two MX Series routers acting as core switches
- One MX Series router acting as an aggregation switch
- Junos OS Release 10.2 or later for MX Series routers

Overview and Topology

Figure 56 on page 866 displays the topology for this example. The topology contains three MX Series routers. CS1 and CS2 act as core routers in the topology, and AS1 acts as an aggregation switch. Each router has two ring instances, ring-1 and ring-2. All nodes on the ring coordinate protection activities by exchanging messages through the Ethernet ring automatic protection switching (R-APS) messaging protocol. Each ring instance has an RPL owner. The ring-1 RPL owner is CS1; the ring-2 RPL owner is CS2. The RPL owners block or unblock the RPL as conditions require and initiate R-APS messages.

Each ring instance has two interface ports (an east interface and a west interface) that participate in the instance. Interface **ge-2/0/8.0**, the west interface on CS2, is the ring protection link end where ring-2's RPL terminates. Interface **ge-3/2/4.0**, the east interface on CS1, is the ring protection link end where ring-1's RPL terminates.

Each ring instance has a data channel. A data channel is a group of bridge domain virtual LAN (VLAN) IDs. All VLAN IDs within the same ring interface share the same data-forwarding properties controlled by the ERP. The data channel on ring-1 is [200, 300]. The data channel on ring-2 is [500, 600].

Two customer site switches are connected to AS1. Customer site 1 uses VLANs 200 and 300. Customer site 2 uses VLANs 500 and 600.

Figure 56: ERP with Multiple Protection Instances Configured on Three MX Series Routers

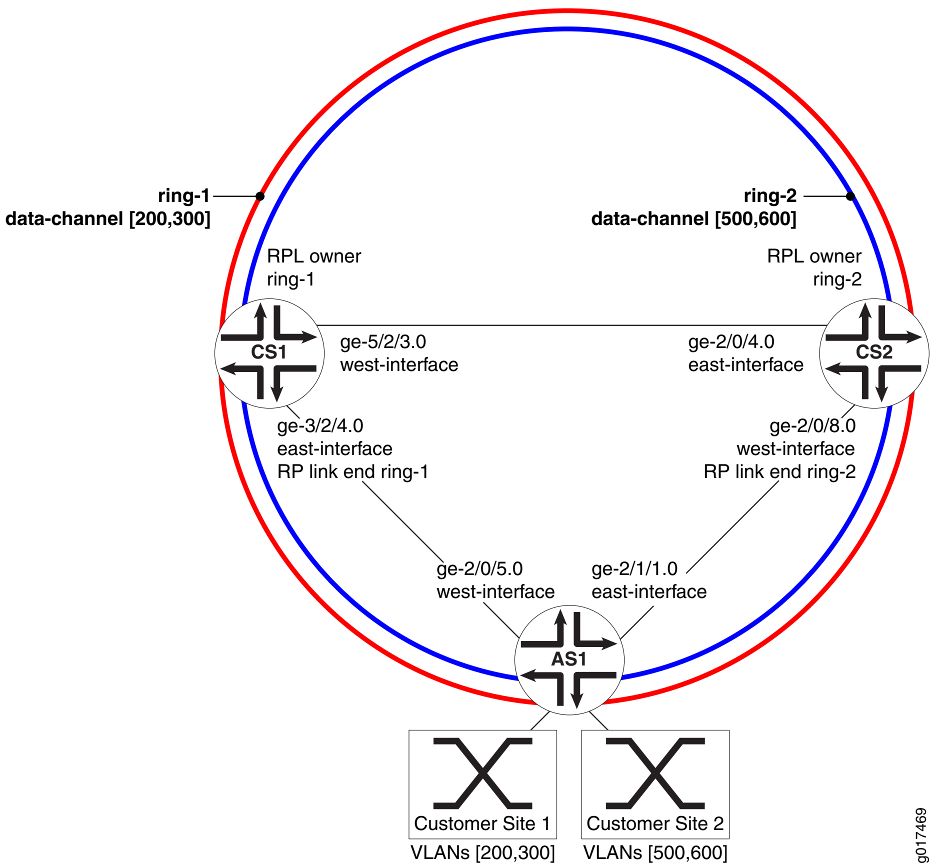

Table 83 on page 866 describes the components of the example topology.

Table 83: Components of the Network Topology

Property	Settings
Ring instances	• **ring-1**—Data channel [200,300] • **ring-2**—Data channel [500,600]
Customer sites	Two customer sites are connected to AS 1: • Customer site 1, VLAN 200 and VLAN 300 • Customer site 2, VLAN 500 and VLAN 600

Table 83: Components of the Network Topology (continued)

Property	Settings

CS1 router	CS1 has the following protection group properties:
	• RPL owner—**ring-1**.
	• East interface—**ge-3/2/4.0**.
	• West interface—**ge-5/2/3.0**.
	• Data channel for **ring-1**—**VLAN 200**, **VLAN 300**.
	• Data channel for **ring-2**—**VLAN 500**, **VLAN 600**.
	• Ring protection link end for **ring-1**—**ge-3/2/4.0**.
	CS1 has the following routing and bridging properties:
	• Routing instance—**vs**.
	• Bridge domains:
	  ◦ **bd100** is associated with **vlan-id 100**.
	  ◦ **bd101** is associated with **vlan-id 101**.
	  ◦ **bd200** is associated with **vlan-id 200**.
	  ◦ **bd300** is associated with **vlan-id 300**.
	  ◦ **bd500** is associated with **vlan-id 500**.
	  ◦ **bd600** is associated with **vlan-id 600**.

CS2 router	CS2 has the following protection group properties: • RPL owner—**ring-2**. • East interface—**ge-2/0/4.0**. • West interface—**ge-2/0/8.0**. • Ring protection link end for **ring-2**—**ge-2/0/8.0**. • Data channel for **ring-1**—**VLAN 200**, **VLAN 300**. • Data channel for **ring-2**—**VLAN 500**, **VLAN 600**. CS2 has the following bridging properties: • **bd100** is associated with **vlan-id 100**. • **bd101** is associated with **vlan-id 101**. • **bd200** is associated with **vlan-id 200**. • **bd300** is associated with **vlan-id 300**. • **bd500** is associated with **vlan-id 500**. • **bd600** is associated with **vlan-id 600**.

Table 83: Components of the Network Topology (continued)

Property	Settings
AS1 router	AS1 has the following protection group properties: • East interface—**ge-2/0/5.0**. • West interface—**ge-2/1/1.0**. • Data channel for **ring-1—VLAN 200**, **VLAN 300**. • Data channel for **ring-2—VLAN 500**, **VLAN 600**. AS1 has the following bridging properties: • **bd100** is associated with **vlan-id 100**. • **bd101** is associated with **vlan-id 101**. • **bd200** is associated with **vlan-id 200**. • **bd300** is associated with **vlan-id 300**. • **bd500** is associated with **vlan-id 500**. • **bd600** is associated with **vlan-id 600**.

Configuration

To enable ERP with ring instances on CS1, CS2, and AS1, perform these tasks:

- Configuring ERP on CS1 on page 868
- Configuring ERP on CS2 on page 871
- Configuring ERP on AS1 on page 874

Configuring ERP on CS1

CLI Quick Configuration

To quickly configure CS1 for ERP, copy the following commands and paste them into the switch terminal window of CS1:

```
[edit]
set interfaces ge-3/2/4 vlan-tagging
set interfaces ge-3/2/4 unit 0 family bridge interface-mode trunk
set interfaces ge-3/2/4 unit 0 family bridge vlan-id-list 100-1000
set interfaces ge-5/2/3 vlan-tagging
set interfaces ge-5/2/3 unit 0 family bridge interface-mode trunk
set interfaces ge-5/2/3 unit 0 family bridge vlan-id-list 100-1000
set protocols protection-group ethernet-ring ring-1 ring-protection-link-owner
set protocols protection-group ethernet-ring ring-1 east-interface control-channel ge-3/2/4.0
set protocols protection-group ethernet-ring ring-1 east-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 east-interface ring-protection-link-end
set protocols protection-group ethernet-ring ring-1 west-interface control-channel ge-5/2/3.0
set protocols protection-group ethernet-ring ring-1 west-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 data-channel vlan [200, 300]
set protocols protection-group ethernet-ring ring-2 east-interface control-channel ge-3/2/4.0
set protocols protection-group ethernet-ring ring-2 east-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 west-interface control-channel ge-5/2/3.0
set protocols protection-group ethernet-ring ring-2 west-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 data-channel vlan [500, 600]
set routing-instances vs instance-type virtual-switch
set routing-instances vs interface ge-3/2/4.0
set routing-instances vs interface ge-5/2/3.0
```



```
set routing-instances vs bridge-domains bd101 vlan-id 101
set routing-instances vs bridge-domains bd200 vlan-id 200
set routing-instances vs bridge-domains bd300 vlan-id 300
set routing-instances vs bridge-domains bd500 vlan-id 500
set routing-instances vs bridge-domains bd600 vlan-id 600
```

Step-by-Step Procedure

To configure ERP on CS1:

1. Configure the trunk interface **ge-3/2/4** to connect CS1 to CS2 and the trunk interface **ge-5/2/3** to connect CS1 to AS, and configure the **family** statement as **bridge** with a VLAN ID list of 100 through 1000:

```
[edit interfaces]
user@cs1# set ge-3/2/4 vlan-tagging
user@cs1# set ge-3/2/4 unit 0 family bridge interface-mode trunk
user@cs1# set ge-3/2/4 unit 0 family bridge vlan-id-list 100-1000
user@cs1# set ge-5/2/3 vlan-tagging
user@cs1# set ge-5/2/3 unit 0 family bridge interface-mode trunk
user@cs1# set ge-5/2/3 unit 0 family bridge vlan-id-list 100-1000
```

2. Enable ERP, specifying the control channels and data channels for **ring-1** and **ring-2**, and configure **ring-1** as the ring protection link owner:

NOTE: Always configure the **east-interface** statement first, before configuring the **west-interface** statement.

```
[edit protection-group]
user@cs1# set ethernet-ring ring-1 ring-protection-link-owner
user@cs1# set ethernet-ring ring-1 east-interface control-channel ge-3/2/4.0
user@cs1# set ethernet-ring ring-1 east-interface control-channel vlan 100
user@cs1# set ethernet-ring ring-1 east-interface ring-protection-link-end
user@cs1# set ethernet-ring ring-1 west-interface control-channel ge-5/2/3.0
user@cs1# set ethernet-ring ring-1 west-interface control-channel vlan 100
user@cs1# set ethernet-ring ring-1 data-channel vlan [200, 300]
user@cs1# set ethernet-ring ring-2 east-interface control-channel ge-3/2/4.0
user@cs1# set ethernet-ring ring-2 east-interface control-channel vlan 101
user@cs1# set ethernet-ring ring-2 west-interface control-channel ge-5/2/3.0
user@cs1# set ethernet-ring ring-2 west-interface control-channel vlan 101
user@cs1# set ethernet-ring ring-2 data-channel vlan [500, 600]
```

3. Configure the routing instance, the bridge domains, and the VLAN IDs associated with each bridge domain:

```
[edit routing-instances]
user@cs1# set vs instance-type virtual-switch
user@cs1# set vs interface ge-3/2/4.0
user@cs1# set vs interface ge-5/2/3.0
user@cs1# set vs bridge-domains bd100 vlan-id 100
user@cs1# set vs bridge-domains bd101 vlan-id 101
user@cs1# set vs bridge-domains bd200 vlan-id 200
user@cs1# set vs bridge-domains bd300 vlan-id 300
user@cs1# set vs bridge-domains bd500 vlan-id 500
```



```
user@cs1# set vs bridge-domains bd600 vlan-id 600
```

Results Check the results of the configuration:

```
user@cs1> show configuration
interfaces {
  ge-3/2/4 {
    vlan-tagging;
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000;
      }
    }
  }
  ge-5/2/3 {
    vlan-tagging;
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000;
      }
    }
  }
  protocols {
    protection-group {
      ethernet-ring ring-1 {
        east-interface {
          control-channel {
            ge-3/2/4.0;
            vlan 100;
          }
          ring-protection-link-end;
        }
        west-interface {
          control-channel {
            ge-5/2/3.0;
            vlan 100;
          }
        }
        data-channel {
          vlan [ 200 300 ];
        }
      }
    }
    protection-group {
      ethernet-ring ring-2{
        east-interface {
          control-channel {
            ge-3/2/4.0;
            vlan 101;
          }
        }
        west-interface {
```



```
            control-channel {
              ge-5/2/3.0;
              vlan 101;
            }
          }
          data-channel {
            vlan [ 500 600 ];
          }
        }
      }
    }
    routing-instances {
      vs {
        instance-type virtual-switch;
        interface ge-3/2/4.0;
        interface ge-5/2/3.0;
        bridge-domains {
          bd100 {
            vlan-id 100;
          }
          bd101 {
            vlan-id 101;
          }
          bd200 {
            vlan-id 200;
          }
          bd300 {
            vlan-id 300;
          }
          bd500 {
            vlan-id 500;
          }
          bd600 {
            vlan-id 600;
          }
        }
      }
    }
  }
```

Configuring ERP on CS2

CLI Quick Configuration

To quickly configure CS2 for ERP, copy the following commands and paste them into the switch terminal window of CS2:

```
[edit]
set interfaces ge-2/0/4 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/4 unit 0 family bridge vlan-id-list 100-1000
set interfaces ge-2/0/8 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/8 unit 0 family bridge vlan-id-list 100-1000
set protocols protection-group ethernet-ring ring-1 east-interface control-channel ge-2/0/4.0
set protocols protection-group ethernet-ring ring-1 east-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 west-interface control-channel ge-2/0/8.0
set protocols protection-group ethernet-ring ring-1 west-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 data-channel vlan [200, 300]
```

```
set protocols protection-group ethernet-ring ring-2 ring-protection-link-owner
set protocols protection-group ethernet-ring ring-2 east-interface control-channel ge-2/0/4.0
set protocols protection-group ethernet-ring ring-2 east-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 west-interface control-channel ge-2/0/8.0
set protocols protection-group ethernet-ring ring-2 west-interface ring-protection-link-end
set protocols protection-group ethernet-ring ring-2 west-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 data-channel vlan [500, 600]
set bridge-domains bd100 vlan-id 100
set bridge-domains bd101 vlan-id 101
set bridge-domains bd200 vlan-id 200
set bridge-domains bd300 vlan-id 300
set bridge-domains bd500 vlan-id 500
set bridge-domains bd600 vlan-id 600
```

Step-by-Step Procedure

To configure ERP on CS2:

1. Configure the trunk interface **ge-2/0/4** to connect CS2 to CS1 and trunk interface **ge-2/0/8** to connect CS2 to CS1, and configure the **family** statement as **bridge** with a VLAN ID list of 100 through 1000:

```
[edit interfaces]
user@cs2# set ge-2/0/4 unit 0 family bridge interface-mode trunk
user@cs2# set ge-2/0/4 unit 0 family bridge vlan-id-list 100-1000
user@cs2# set ge-2/0/8 unit 0 family bridge interface-mode trunk
user@cs2# set ge-2/0/8 unit 0 family bridge vlan-id-list 100–1000
```

2. Enable ERP, specifying the control channels and data channels for **ring-1** and **ring-2**, and configure **ring-2** as the ring protection link owner:

NOTE: Always configure the **east-interface** statement first, before configuring the **west-interface** statement.

```
[edit protection-group]
user@cs2# set ethernet-ring ring-1 east-interface control-channel ge-2/0/4.0
user@cs2# set ethernet-ring ring-1 east-interface control-channel vlan 100
user@cs2# set ethernet-ring ring-1 west-interface control-channel ge-2/0/8.0
user@cs2# set ethernet-ring ring-1 west-interface control-channel vlan 100
user@cs2# set ethernet-ring ring-2 data-channel vlan [200, 300]
user@cs2# set ethernet-ring ring-2 east-interface control-channel ge-2/0/4.0
user@cs2# set ethernet-ring ring-2 east-interface control-channel vlan 101
user@cs2# set ethernet-ring ring-2 ring-protection-link-owner
user@cs2# set ethernet-ring ring-2 west-interface control-channel ge-2/0/8.0
user@cs2# set ethernet-ring ring-2 west-interface control-channel vlan 101
user@cs2# set ethernet-ring ring-2 west-interface ring-protection-link-end
user@cs2# set ethernet-ring ring-2 data-channel vlan [500, 600]
```

3. Configure the routing instance, the bridge domains, and the VLAN IDs associated with each bridge domain:

```
[edit bridge-domains]
user@cs2# set bd100 vlan-id 100
user@cs2# set bd101 vlan-id 101
user@cs2# set bd200 vlan-id 200
```

```
user@cs2# set bd300 vlan-id 300
user@cs2# set bd500 vlan-id 500
user@cs2# set bd600 vlan-id 600
```

Results Check the results of the configuration:

```
user@cs2> show configuration
interfaces {
  ge-2/0/4 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000;
      }
    }
  }
  ge-2/0/8 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000;
      }
    }
  }
  protocols {
    protection-group {
      ethernet-ring ring-1 {
        east-interface {
          control-channel {
            ge-2/0/4.0;
            vlan 100;
          }
        }
        west-interface {
          control-channel {
            ge-2/0/8.0;
            vlan 100;
          }
        }
        data-channel {
          vlan [200, 300];
        }
      }
    }
    ethernet-ring ring-2 {
      east-interface {
        control-channel {
          ge-2/0/4.0;
          vlan 101;
        }
      }
      west-interface {
        control-channel {
          ge-2/0/8.0;
```

```
                vlan 101;
              }
              ring-protection-link-end;
            }
            data-channel {
              vlan [500, 500];
            }
          }
        }
        bridge-domains {
          bd100 {
            vlan-id 100;
          }
          bd101 {
            vlan-id 101;
          }
          bd200 {
            vlan-id 200;
          }
          bd300 {
            vlan-id 300;
          }
          bd500 {
            vlan-id 500;
          }
          bd600 {
            vlan-id 600;
          }
        }
      }
```

Configuring ERP on AS1

CLI Quick Configuration

To quickly configure AS1 for ERP, copy the following commands and paste them into the switch terminal window of AS1:

```
[edit]
set interfaces ge-2/0/5 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/5 unit 0 family bridge vlan-id-list 100-1000
set interfaces ge-2/1/1 unit 0 family bridge interface-mode trunk
set interfaces ge-2/1/1 unit 0 family bridge vlan-id-list 100-1000
set protocols protection-group ethernet-ring ring-1 east-interface control-channel ge-2/0/5.0
set protocols protection-group ethernet-ring ring-1 east-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 west-interface control-channel ge-2/1/1.0
set protocols protection-group ethernet-ring ring-1 west-interface control-channel vlan 100
set protocols protection-group ethernet-ring ring-1 data-channel vlan [200, 300]
set protocols protection-group ethernet-ring ring-2 east-interface control-channel ge-2/0/5.0
set protocols protection-group ethernet-ring ring-2 east-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 west-interface control-channel ge-2/1/1.0
set protocols protection-group ethernet-ring ring-2 west-interface control-channel vlan 101
set protocols protection-group ethernet-ring ring-2 data-channel vlan [500, 600]
set bridge-domains bd100 vlan-id 100
set bridge-domains bd101 vlan-id 101
set bridge-domains bd200 vlan-id 200
set bridge-domains bd300 vlan-id 300
set bridge-domains bd500 vlan-id 500
```

set bridge-domains bd600 vlan-id 600

Step-by-Step Procedure

To configure ERP on AS1:

1. Configure the trunk interface **ge-2/0/5** to connect CS2 to CS1 and trunk interface **ge-2/1/1** to connect CS2 to CS1, and configure the **family** statement as **bridge** with a VLAN ID list of 100 through 1000:

```
[edit interfaces]
user@as1# set ge-2/0/5 unit 0 family bridge interface-mode trunk
user@as1# set ge-2/0/5 unit 0 family bridge vlan-id-list 100-1000
user@as1# set ge-2/1/1 unit 0 family bridge interface-mode trunk
user@as1# set ge-2/1/1 unit 0 family bridge vlan-id-list 100
```

2. Enable ERP, specifying the control channels and data channels for **ring-1** and **ring-2**:

NOTE: Always configure the **east-interface** statement first, before configuring the **west-interface** statement.

```
[edit protection-group]
user@as1# set ethernet-ring ring-1 east-interface control-channel ge-2/0/5.0
user@as1# set ethernet-ring ring-1 east-interface control-channel vlan 100
user@as1# set ethernet-ring ring-1 west-interface control-channel ge-2/1/1.0
user@as1# set ethernet-ring ring-1 west-interface control-channel vlan 100
user@as1# set ethernet-ring ring-2 east-interface control-channel ge-2/0/5.
user@as1# set ethernet-ring ring-2 east-interface control-channel vlan 101
user@as1# set ethernet-ring ring-2 west-interface control-channel ge-2/1/1.0
user@as1# set ethernet-ring ring-2 west-interface control-channel vlan 101
user@as1# set ethernet-ring ring-2 data-channel vlan [500, 600]
```

3. Configure the routing instance, the bridge domains, and the VLAN IDs associated with each bridge domain:

```
[edit bridge-domains]
user@as1# set bd100 vlan-id 100
user@as1# set bd101 vlan-id 101
user@as1# set bd200 vlan-id 200
user@as1# set bd300 vlan-id 300
user@as1# set bd500 vlan-id 500
user@as1# set bd600 vlan-id 600
```

Results

Check the results of the configuration:

```
user@as1> show configuration
interfaces {
  ge-2/0/5 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 100-1000
      }
```

```
        }
    }
    ge-2/1/1 {
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list 100-1000
            }
        }
    }
    protocols {
        protection-group {
            ethernet-ring ring-1 {
                east-interface {
                    control-channel {
                        ge-2/0/5.0;
                        vlan 100;
                    }
                }
                west-interface {
                    control-channel {
                        ge-2/1/1.0;
                        vlan 100;
                    }
                }
                data-channel {
                    vlan [200, 300];
                }
            }
        }
    }
    protection-group {
        ethernet-ring ring-2 {
            east-interface {
                control-channel {
                    ge-2/0/5.0;
                    vlan 101;
                }
            }
            west-interface {
                control-channel {
                    ge-2/1/1.0;
                    vlan 101;
                }
            }
            data-channel {
                vlan [500, 600];
            }
        }
    }
    bridge-domains {
        bd100 {
            vlan-id 100;
        }
        bd101 {
            vlan-id 101;
```



```
        }
        bd200 {
          vlan-id 200;
        }
        bd300 {
          vlan-id 300;
        }
        bd500 {
          vlan-id 500;
        }
        bd600 {
          vlan-id 600;
        }
      }
    }
```

Verification

To confirm that the ERP configuration for multiple ring instances is operating, perform these tasks:

Verifying the Ethernet Protection Ring on CS1

Purpose Verify that ERP is enabled on CS1.

Action Show the status of the ring automatic protection switching (R-APS) messages to determine if there is a ring failure:

```
user@cs1>show protection-group ethernet-ring aps

Ethernet Ring Name   Request/state  No Flush  Ring Protection  Originator  Remote
 Node ID
                                               Link Blocked
ring-1               NR             No         Yes              Yes
ring-2               NR             No         Yes              No
13:22:af:31:fc:00
```

Meaning The output displayed shows that protection groups **ring-1** and **ring-2** have a **Request/state** of **NR**, meaning there is no request for APS on the ring. If a **Request/state** of **SF** is displayed, it indicates there is a signal failure on the ring. The output also shows that the ring protection link is not blocked. The **No Flush** field displays **No**, indicating that MAC addresses will be flushed when the ring nodes receive this message first time. A value of **Yes** would indicate MAC address flushing is not needed. The **Originator** field for **ring-1** displays **yes**, indicating that this node is an R-APS originator.

Verifying the Data Channel CS1

Purpose Verify the forwarding state of the data channel.

Action List the interfaces acting as the control channels and their respective data channels (represented by the Spanning Tree Protocol (STP) index number):

```
user@cs1>show protection-group ethernet-ring data-channel
Ethernet ring data channel parameters for protection group ring-1

Interface     STP index   Forward State
ge-3/2/4      122         forwarding
ge-5/2/3      123         forwarding


Ethernet ring data channel parameters for protection group ring-2

Interface     STP index   Forward State

ge-3/2/4      124         discarding
ge-5/2/3      125         forwarding
```

Meaning The output displayed shows the STP index number used by each interface in ring instances **ring-1** and **ring-2**. The STP index controls the forwarding behavior for a set of VLANs on the data channel of a ring instance on a ring interface. For ring instances, there are multiple STP index numbers (here representing VLANs 200, 300, 500, and 600). The **Forward State** shows whether the data channel is **forwarding** or **discarding** traffic.

Verifying the VLANs on CS1

Purpose Verify the data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Action List dynamic VLAN membership:

```
user@cs1>show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group ring-1

Interface     Vlan     STP Index    Bridge Domain
ge-3/2/4      200      122          vs/bd200
ge-5/2/3      200      123          vs/bd200
ge-3/2/4      300      122          vs/bd300
ge-5/2/3      300      123          vs/bd300

Ethernet ring IFBD parameters for protection group ring-2
```



```
Interface     Vlan     STP Index     Bridge Domain
ge-3/2/4      500      124           vs/bd500
ge-5/2/3      500      125           vs/bd500
ge-3/2/4      600      124           vs/bd600
ge-5/2/3      600      125           vs/bd600
```

Meaning The output displayed shows the ring interfaces **ge-3/2/4** and **ge-5/2/3** in protection groups **ring-1** and **ring-2**. For **ring-1**, VLAN 200 and VLAN 300 are being supported on both **STP Index 122** and **123** on bridge domains **bd200** and **bd300**. For **ring-2**, VLAN 500 and VLAN 600 are being supported on both **STP Index 124** and **125** on bridge domains **bd500** and **bd600**. The data channel controls the traffic on the VLAN IDs to facilitate load balancing.

Verifying the Ethernet Protection Ring on CS2

Purpose Verify that ERP is enabled on CS2.

Action Show the status of the ring APS (R-APS) messages to determine if there is a ring failure:

```
user@cs2>show protection-group ethernet-ring aps

Ethernet Ring Name    Request/state  No Flush  Ring Protection  Originator  Remote
 Node ID
                                               Link Blocked
Ring-1                 NR             No        No                No
00:21:59:03:ff:d0
Ring-2                 NR             No        Yes               Yes
```

Meaning The output displayed shows that protection groups **ring-1** and **ring-2** have a **Request/state** of **NR**, meaning there is no request for APS on the ring. If a **Request/state** of **SF** is displayed, it indicates there is a signal failure on the ring. The output also shows that the ring protection link is not blocked. The **No Flush** field displays **No**, indicating that MAC addresses will be flushed when the ring nodes receive this message first time. A value of **Yes** would indicate MAC address flushing is not needed. The **Originator** field for **ring-1** displays **yes**, indicating that this node is an R-APS originator. The **Originator** field for **ring-2** displays **No**, indicating that this node is not an R-APS originator.

Verifying the Data Channel CS2

Purpose Verify the forwarding state of the data channel.

Action List the interfaces acting as the control channels and their respective data channels (represented by the STP index number):

```
user@cs2> show protection-group ethernet-ring data-channel
Ethernet ring data channel parameters for protection group ring-1
```



```
Interface     STP index  Forward State
ge-2/0/4      44         forwarding
ge-2/0/8      45         forwarding


Ethernet ring data channel parameters for protection group ring-2

Interface     STP index  Forward State
ge-2/0/4      46         forwarding
ge-2/0/8      47         discarding
```

Meaning The output displayed shows the STP index number used by each interface in ring instances **ring-1** and **ring-2**. The STP index controls the forwarding behavior for a set of VLANs on the data channel of a ring instance on a ring interface. For ring instances, there are multiple STP index numbers (here representing VLANs 200, 300, 500, and 600). The **Forward State** shows whether the data channel is **forwarding** or **discarding** traffic.

Verifying the VLANs on CS2

Purpose Verify the data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Action List dynamic VLAN membership:

```
user@cs2> show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group ring-1

Interface     Vlan     STP Index     Bridge Domain
ge-2/0/4      200      44            default-switch/bd200
ge-2/0/8      200      45            default-switch/bd200
ge-2/0/4      300      44            default-switch/bd300
ge-2/0/8      300      45            default-switch/bd300

Ethernet ring IFBD parameters for protection group ring-2

Interface     Vlan     STP Index     Bridge Domain
ge-2/0/4      500      46            default-switch/bd500
ge-2/0/8      500      47            default-switch/bd500
ge-2/0/4      600      46            default-switch/bd600
ge-2/0/8      600      47            default-switch/bd600
```

Meaning The output displayed shows the ring interfaces **ge-2/0/4** and **ge-2/0/8** in protection groups **ring-1** and **ring-2**. For **ring-1**, VLAN 200 and VLAN 300 are being supported on both **STP Index 44** and **45** on bridge domains **bd200** and **bd300**. For **ring-2**, VLAN 500 and VLAN 600 are being supported on both **STP Index 46** and **47** on bridge domains **bd500** and **bd600**. The data channel controls the traffic on the VLAN IDs to facilitate load balancing.

Verifying the Ethernet Protection Ring on AS1

Purpose Verify that ERP is enabled on AS1.

Action Show the status of the ring APS (R-APS) messages to determine if there is a ring failure:

```
user@as1> show protection-group ethernet-ring aps

Ethernet Ring Name    Request/state  No Flush  Ring Protection  Originator  Remote
 Node ID
                                                 Link Blocked
Ring-1                 NR             No         Yes               No
00:21:59:03:ff:d0
Ring-2                 NR             No         Yes               No
13:22:af:31:fc:00
```

Meaning The output displayed shows that protection groups **ring-1** and **ring-2** have a **Request/state** of **NR**, meaning there is no request for APS on the ring. If a **Request/state** of **SF** is displayed, it indicates there is a signal failure on the ring. The output also shows that the ring protection link is not blocked. The **No Flush** field displays **No**, indicating that MAC addresses will be flushed when the ring nodes receive this message first time. A value of **Yes** would indicate MAC address flushing is not needed. The **Originator** field for **ring-1** and **ring-2** displays **No**, indicating that this node is not the R-APS originator.

Verifying the Data Channels on AS1

Purpose Verify the forwarding state of the data channel.

Action List the interfaces acting as the control channels and their respective data channels (represented by the STP index number):

```
user@as1>show protection-group ethernet-ring data-channel
Ethernet ring data channel parameters for protection group ring-1

Interface     STP index  Forward State
ge-2/0/5      22         forwarding
ge-2/1/1      23         forwarding


Ethernet ring data channel parameters for protection group ring-2

Interface     STP index  Forward State
ge-2/0/5      24         forwarding
ge-2/1/1      25         forwarding
```

Meaning The output displayed shows the STP index number used by each interface in ring instances **ring-1** and **ring-2**. The STP index controls the forwarding behavior for a set of VLANs on the data channel of a ring instance on a ring interface. For ring instances, there are multiple STP index numbers (here representing VLANs 200, 300, 500, and 600). The **Forward**

State shows whether the data channel is **forwarding** or **discarding** traffic. All data channels are forwarding traffic.

Verifying the VLANs on AS1

Purpose Verify the data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Action List dynamic VLAN membership:

```
user@as1>show protection-group ethernet-ring vlan

Ethernet ring IFBD parameters for protection group ring-1

Interface     Vlan     STP Index     Bridge Domain
ge-2/0/5      200      22            default-switch/bd200
ge-2/1/1      200      23            default-switch/bd200
ge-2/0/5      300      22            default-switch/bd300
ge-2/1/1      300      23            default-switch/bd300

Ethernet ring IFBD parameters for protection group ring-2

Interface     Vlan     STP Index     Bridge Domain
ge-2/0/5      500      24            default-switch/bd500
ge-2/1/1      500      25            default-switch/bd500
ge-2/0/5      600      24            default-switch/bd600
ge-2/1/1      600      25            default-switch/bd600
```

Meaning The output displayed shows the ring interfaces **ge-2/0/5** and **ge-2/1/1** in protection groups **ring-1** and **ring-2**. For **ring-1**, VLAN 200 and VLAN 300 are being supported on both **STP Index 22** and **23** on bridge domains **bd200** and **bd300**. For **ring-2**, VLAN 500 and VLAN 600 are being supported on both **STP Index 24** and **25** on bridge domains **bd500** and **bd600**. The data channel controls the traffic on the VLAN IDs to facilitate load-balancing.

Related Documentation

- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857
- Ethernet Ring Protection on page 855

Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation

Under normal operating conditions, when Ethernet ring protection is configured correctly, the ring protection link (RPL) owner (Router 1 in the configuration example) will see the following:

Router 1 Operational Commands (Normal Ring Operation)

```
user@router1> show protection-group ethernet-ring aps
Ethernet Ring Name   Request/state   No Flush   Ring Protection Link Blocked
pg101                NR              No         Yes
```



```
Originator  Remote Node ID
Yes
```

Note that the ring protection link is blocked and the node is marked as the originator of the protection.

```
user@router1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/0/1     ge-1/0/1.1       discarding     Yes
ge-1/2/4     ge-1/2/4.1       forwarding     No

Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```

Note that the protection interface is discarding while the other interface is forwarding.

```
user@router1> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         Ring Protection Link Owner
pg101            idle         NR-RB         Yes

Restore Timer  Quard Timer  Operation state
disabled       disabled    operational
```

Note that Router 1 is the owner and timers are disabled.

```
user@router1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                         : 1
RAPS received                     : 0
Local SF happened:                : 0
Remote SF happened:               : 0
NR event happened:                : 0
NR-RB event happened:             : 1
```

Note that only minimal RAPS messages have been sent to establish the ring.

Under normal operating conditions, the other routers on the ring (Router 2 and Router 3) will see the following similar output:

Router 2 and Router 3 Operational Commands (Normal Ring Operation)

```
user@router2> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               NR             No        Yes

Originator  Remote Node ID
No          00:01:01:00:00:01
```

Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102
```

```
Interface     Control Channel  Forward State  Ring Protection Link End
ge-1/2/1      ge-1/2/1.1       forwarding     No
ge-1/0/2      ge-1/0/2.1       forwarding     No

Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```

Note that both interfaces are forwarding. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event        Ring Protection Link Owner
pg102             idle          NR-RB        No

Restore Timer  Quard Timer  Operation state
disabled       disabled    operational
```

Note that Router 2 is not the owner. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                         : 0
RAPS received                     : 1
Local SF happened:                : 0
Remote SF happened:               : 0
NR event happened:                : 0
NR-RB event happened:             : 1
```

Router 3 will see almost identical information.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Ring Protection on page 855
- Example: Configuring Ethernet Ring Protection for MX Series Routers on page 858
- Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition on page 884

Example: Viewing Ethernet Ring Protection Status—Ring Failure Condition

This section assumes that Ethernet ring protection is configuring correctly, that Router 1 is the ring protection link (RPL) owner, and that there is a link failure between Router 2 and Router 3 in the configuration example.

Router 1 Operational Commands (Ring Failure Condition)

```
user@router1> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               SF             NO        No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

Note that the ring protection link is no longer blocked and the node is no longer marked as originator.

```
user@router1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/0/1     ge-1/0/1.1       forwarding     Yes
ge-1/2/4     ge-1/2/4.1       forwarding     No

Signal Failure  Admin State
Clear           IFF ready
Clear           IFF ready
```

Note that the protection interface is now forwarding (so is the other interface).

```
user@router1> show protection-group ethernet-ring node-state
how protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         Ring Protection Link Owner
pg101            protected    SF            Yes

Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

Note that Router 1 has recorded the span failure (SF).

```
user@router1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                        : 1
RAPS received                    : 1
Local SF happened:               : 0
Remote SF happened:              : 1
NR event happened:               : 0
NR-RB event happened:            : 1
```

Note that the R-APS messages have recorded the remote failure.

Under a failure condition, the other routers on the ring (Router 2 and Router 3) will see the following similar output:

Router 2 and Router 3 Operational Commands (Failure Condition)

```
user@router2> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               SF             No        No

Originator  Remote Node ID
Yes         00:00:00:00:00:00
```

Note the failure event (SF). Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/2/1     ge-1/2/1.1       forwarding     No
```

```
ge-1/0/2      ge-1/0/2.1       discarding     No

Signal Failure  Admin State
Clear           IFF ready
set             IFF ready
```

Note that the failed interface (**ge-1/0/2.1**) is not forwarding. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         Ring Protection Link Owner
pg102             idle          NR-RB         No

Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

Note that Router 2 is not the owner. Router 3 will see almost identical information.

```
user@router2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                       : 1
RAPS received                   : 1
Local SF happened:              : 1
Remote SF happened:             : 0
NR event happened:              : 0
NR-RB event happened:           : 1
```

Note that the R-APS messages have recorded the remote failure. Router 3 will see almost identical information.

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Ring Protection on page 855
- Example: Configuring Ethernet Ring Protection for MX Series Routers on page 858
- Example: Viewing Ethernet Ring Protection Status—Normal Ring Operation on page 882

CHAPTER 35

CFM Action Profile to Bring Down a Group of Logical Interfaces

CFM Action Profile to Bring Down a Group of Logical Interfaces Overview

With growing networks, there is a requirement of monitoring a large number of services using CFM. To monitor each service, one session per service logical interface is required. If the services are large in number, this method does not scale as the number of sessions are limited. Instead of one CFM session per service, a single CFM session can monitor multiple services.

Also, there are scenarios where the user-to-network interface (UNI) device needs to be brought down based on sessions on network-to-network Interface (NNI) logical interface. Here, the NNI logical interface refers to core interface and UNI physical interface refers to access interface hosting multiple service logical interfaces. Based on core interface monitoring, you can bring down service logical interfaces associated with access interface.

Figure 57 on page 888 illustrates a topology where a number of services destined to customer-edge (CE) routers share a single port on a provider-edge (PE) router. Each service uses one logical interface. A set of services or logical interfaces (colored in yellow) are destined to one CE router and a set of services or logical interfaces colored in red are destined to another CE router. To monitor each service, you need dedicated down maintenance association end point (MEP) sessions for each service. You can bring down the service by bringing down the service logical interface whenever the session goes down. However, this approach is not scalable if we have large number of services. Monitoring the CFM session on the physical interface is also not feasible because multiple CE routers might be connected and the services to other CE router could be disrupted. To address this issue of monitoring multiple services with a single session, you can create a CCM action profile to bring down a group of logical interfaces by using a CFM session that is configured on a single logical interface.

Figure 57: Topology of Multiple VLAN Services Sharing a Single Port on PE Router Destined to Multiple CE Routers

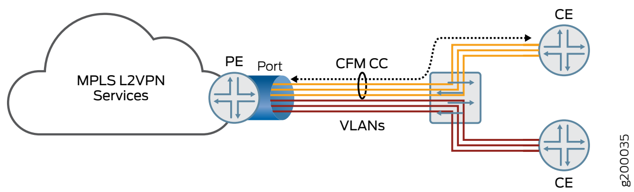

You can configure CCM action profiles for the following scenarios:

- To bring down a group of logical interfaces all having the same parent port when CCM monitoring session is running on one of the logical interface but on a different parent port.
- To bring down a group of logical interfaces when CCM monitoring session is running on one of the logical interfaces, all belonging to the same parent port.
- To bring down the port, when the CCM monitoring session is running on one of the logical interfaces of a different parent port.

Benefits of Creating CFM Action Profile to Bring Down a Group of Logical Interfaces

- Reduces resource requirement in scaled networks where multiple services need to be monitored.
- Avoids the need to create individual MEP sessions for each service in a topology that includes multiple services to be monitored, thereby enhancing the performance and scalability of the network.

Related Documentation

- Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces on page 888

Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces

To monitor multiple services or IFLs using CFM session configured on a single logical interface, you can create a CCM action profile to bring down a group of logical interfaces. You need to define an action to bring down the interface group in the action profile. You will then define the interface device name and the number of logical interfaces that have to be brought down. A logical interface is represented by a combination of the **interface-device-name** and **unit-list**. The following steps explain the procedure to bring down a group of logical interfaces when the **interface-device-name** and/or **unit-list** are specified.

1. In configuration mode, at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level, specify the name of the action profile and the CFM event(s). You can configure more than one event in the action profile.

   [edit protocols oam ethernet connectivity-fault-management]
   user@host# set action-profile *profile-name* event [*event1*, *event2*, *event3*..]

For example,

```
user@host# set action-profile AP_test event adjacency-loss rdi
```

NOTE: The action **interface-group-down** will not be supported with events other than adjacency-loss and RDI. Any other events configured results in a commit error.

2. In configuration mode, at the [**edit protocols oam ethernet connectivity-fault-management action-profile *profile-name***] hierarchy level, define the action to bring down the interface group.

```
[edit protocols oam ethernet connectivity-fault-management action-profile AP-test ]
user@host# set action interface-group-down
```

NOTE: The action **interface-group-down** will not be supported with other interface related actions. Any other actions configured results in a commit error.

3. At the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level, define the maintenance domain. Specify the maintenance-association parameters.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set maintenance-domain domain-name  level number
  maintenance-association ma-name continuity-check interval 1s
```

For example,

```
user@host# set maintenance-domain md6 level 6 maintenance-association ma6
  continuity-check interval 1s
```

4. At the **edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name***, define the maintenance association endpoint and the associated parameters.

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

For example,

```
user@host# set mep 101 interface ge-0/0/0.0 direction down remote -mep 102
```

5. If the action-profile has **interface-group-down** action configured, it is mandatory to configure the **interface-group** at the RMEP level. In the configuration mode at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id* action-profile *profile-name*** include the **interface-group** statement to bring down the interface group marked with the action profile as **interface-group-down**.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name mep mep-id remote-mep mep-id
  action-profile profile-name]
user@host# set interface-group
```

For example,

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md6 maintenance-association ma6 mep 101 remote-mep 102 action-profile AP_test]
user@host# set interface-group
```

NOTE: If the **interface-group** configuration is not included in the RMEP configuration. The configuration results in commit error.

6. A logical interface is represented by a combination of the **interface-device-name** and **unit-list**. Configure the device interface name and the number of logical interfaces at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id* action-profile *profile-name* interface-group**.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name mep mep-id remote-mep mep-id
  action-profile profile-name interface-group]
user@host# set interface interface-name
user@host# set unit-list logical-interface-unit-number
```

For example,

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md6 maintenance-association ma6 mep 101 remote-mep 102 action-profile AP_test
  interface-group]
user@host# set interface ge-0/0/0.0
user@host# set unit-list 1223-3344
```

In this configuration example, the interface ge-0/0/0.0 is brought down.

NOTE:

- At least one of the **interface-group** parameters, **interface-device-name** or **unit-list** must be configured. If the interface device name is not configured, the MEP interface is considered as the device name and the logical interface on that device is brought down.
- If the **unit-list** parameter exceeds the recommended limit, a commit error occurs.
- If the *interface-device-name* is not specified in the **interface-group**, the logical interface numbers mentioned in **unit-list** for the physical interface is brought down.
- If the *unit-list* is not specified in the **interface-group**, IFLs are brought down for the configured interface.

7. Verify the configuration using **show protocols oam** command.

   [edit]
   user@host# **show protocols oam**

```
ethernet {
    connectivity-fault-management {
        action-profile AP_TEST {
            event {
                adjacency-loss;
                rdi;
            }
            action {
                interface-group-down;
            }
        }
        maintenance-domain md6 {
            level 6;
            maintenance-association ma6 {
                continuity-check {
                    interval 1s;
                }
                mep 102 {
                    interface ge-0/0/0.0;
                    direction down;
                    remote-mep 103 {
                        action-profile AP_TEST;
                        interface-group {
                          ge-0/0/1;
                           unit-list [12 23-33 44];
                        }
                    }
                }
            }
        }
    }
}
```

Related Documentation

- CFM Action Profile to Bring Down a Group of Logical Interfaces Overview on page 887

PART 4

Troubleshooting Information

CHAPTER 36

Monitoring and Troubleshooting Ethernet Interfaces

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, *N*xDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own Data Link Layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own Data Link Layer packets, you can assume that the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

The following types of loopback testing are supported by Junos OS:

- DCE local—Loops packets back on the local data circuit-terminating equipment (DCE).
- DCE remote—Loops packets back on the remote DCE.
- Local—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- Payload—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- Remote—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 84 on page 896 shows the loopback modes supported on the various interface types.

Table 84: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	“Configuring Ethernet Loopback Capability” on page 18
Circuit Emulation E1	Local and remote	*Configuring E1 Loopback Capability*
Circuit Emulation T1	Local and remote	*Configuring T1 Loopback Capability*
E1 and E3	Local and remote	*Configuring E1 Loopback Capability* and *Configuring E3 Loopback Capability*
*N*xDS0	Payload	*Configuring NxDS0 IQ and IQE Interfaces*, *Configuring T1 and NxDS0 Interfaces*, *Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)*, *Configuring Fractional E1 IQ and IQE Interfaces*, and *Configuring Channelized T3 IQ Interfaces*

Table 84: Loopback Modes by Interface Type (continued)

Interface	Loopback Modes	Usage Guidelines
Serial (V.35 and X.21)	Local and remote	*Configuring Serial Loopback Capability*
Serial (EIA-530)	DCE local, DCE remote, local, and remote	*Configuring Serial Loopback Capability*
SONET/SDH	Local and remote	*Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External*
T1 and T3	Local, payload, and remote	*Configuring T1 Loopback Capability* and *Configuring T3 Loopback Capability* See also *Configuring the T1 Remote Loopback Response*

To configure loopback testing, include the **loopback** statement:

user@host# loopback *mode*;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]
- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

Configuring BERT Testing

To configure BERT:

- Configure the duration of the test.

  [edit interfaces *interface-name interface-type*-options]
  user@host#bert-period *seconds*;

  You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs. By default, the BERT period is 10 seconds.

- Configure the error rate to monitor when the inbound pattern is received.

[edit interfaces *interface-name interface-type*-options]
user@host#bert-error-rate *rate*;

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

- Configure the bit pattern to send on the transmit path.

  [edit interfaces *interface-name interface-type*-options]
  user@host#bert-algorithm *algorithm*;

  algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a **?** after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]

user@host# set bert-algorithm ?
Possible completions:
  pseudo-2e11-o152     Pattern is 2^11 -1 (per 0.152 standard)
  pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.152 standard)
  pseudo-2e20-o151     Pattern is 2^20 - 1 (per 0.151 standard)
  pseudo-2e20-o153     Pattern is 2^20 - 1 (per 0.153 standard)
        ...
```

For specific hierarchy information, see the individual interface types.

NOTE: The four-port E1 PIC supports only the following algorithms:

```
pseudo-2e11-o152     Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151     Pattern is 2^23 (per 0.151 standard)
```

When you issue the **help** command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e11-o152      Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7            Pattern is 2^7 - 1
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4      1 bit in 4 is set
repeating-1-in-8      1 bit in 8 is set
repeating-3-in-24     3 bits in 24 are set
```

When you issue the **help** command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e9-o153      Pattern is 2^9  -1 (per 0.153 (511 type) standard)
pseudo-2e11-o152     Pattern is 2^11 -1 (per 0.152 and 0.153 (2047 type)
 standards)
pseudo-2e15-o151      Pattern is 2^15 -1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 -1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 -1 (per 0.153 standard)
pseudo-2e23-o151      Pattern is 2^23 -1 (per 0.151 standard)
repeating-1-in-4      1 bit in 4 is set
repeating-1-in-8      1 bit in 8 is set
repeating-3-in-24     3 bits in 24 are set
```

When you issue the **help** command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

```
all-ones-repeating            Repeating one bits
all-zeros-repeating           Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros        Alternating ones and zeros
repeating-1-in-4              1 bit in 4 is set
repeating-1-in-8              1 bit in 8 is set
repeating-3-in-24             3 bits in 24 are set
pseudo-2e9-o153               Pattern is 2^9 - 1 (per 0.153 standard)
pseudo-2e11-o152              Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151              Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151              Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153              Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151              Pattern is 2^23 (per 0.151 standard)
```

Table 85 on page 900 shows the BERT capabilities for various interface types.

Table 85: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)	—	• Limited algorithms
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)	—	• Limited algorithms
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	• Single port at a time • Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	• Single port at a time
Channelized OC12	—	Yes (channel 0–11)	• Single channel at a time • Limited algorithms • No bit count
Channelized STM1	Yes (channel 0–62)	—	• Multiple channels • Only one algorithm • No error insert • No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	• Multiple ports and channels • Limited algorithms for T1 • No error insert for T1 • No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see *Channelized IQ and IQE Interfaces Properties*.

Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
disable;
```

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name* *interface-type*-bert-start** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you want to terminate the test sooner, issue the **test interface *interface-name* *interface-type*-bert-stop** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the CLI Explorer.

NOTE: To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

Related Documentation

- show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) on page 1803

Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Purpose To monitor Fast Ethernet and Gigabit Ethernet interfaces and begin the process of isolating interface problems when they occur.

Action Table 86 on page 902 provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

Table 86: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
"Monitor Fast Ethernet and Gigabit Ethernet Interfaces" on page 903	
1. Display the Status of Fast Ethernet Interfaces on page 903	**show interfaces terse (fe* | ge*)**
2. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 905	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
3. Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface on page 907	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
4. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 910	**monitor interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
5. Fiber-Optic Ethernet Interface Specifications on page 911	

Meaning You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

Related Documentation

- Display the Status of Gigabit Ethernet Interfaces on page 904
- Display the Status of Fast Ethernet Interfaces on page 903

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Purpose To monitor Fast Ethernet and Gigabit Ethernet interfaces and begin the process of isolating interface problems when they occur.

Action Table 86 on page 902 provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

Table 87: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
"Monitor Fast Ethernet and Gigabit Ethernet Interfaces" on page 903	
1. Display the Status of Fast Ethernet Interfaces on page 903	**show interfaces terse (fe* | ge*)**
2. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 905	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**

Table 87: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces (continued)

Tasks	Command or Action
3. Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface on page 907	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
4. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 910	**monitor interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
5. Fiber-Optic Ethernet Interface Specifications on page 911	

Meaning You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

See Also

- Display the Status of Gigabit Ethernet Interfaces on page 904
- Display the Status of Fast Ethernet Interfaces on page 903

Monitor Fast Ethernet and Gigabit Ethernet Interfaces

By monitoring Fast Ethernet and Gigabit Ethernet interfaces, you begin to isolate Fast Ethernet and Gigabit Ethernet interface problems when they occur.

To monitor your Fast Ethernet and Gigabit Ethernet interfaces, follow these steps:

1. Display the Status of Fast Ethernet Interfaces on page 903
2. Display the Status of Gigabit Ethernet Interfaces on page 904
3. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 905
4. Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface on page 907
5. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 910

Display the Status of Fast Ethernet Interfaces

Purpose To display the status of Fast Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Action user@host> **show interfaces terse (fe* | ge*)**

Sample Output

```
user@host> show interfaces terse fe*
Interface       Admin Link Proto Local                 Remote
fe-2/1/0        up    up
fe-2/1/0.0      up    up   inet  10.116.115.217/29
fe-3/0/2        up    down
fe-3/0/2.0      up    down
```



```
fe-3/0/3        up    up
fe-3/0/3.0      up    up   inet  192.168.223.65/30
fe-4/1/0        down  up
fe-4/1/0.0      up    down inet  10.150.59.133/30
fe-4/1/1        up    up
fe-4/1/1.0      up    up   inet  10.150.59.129/30
fe-4/1/2        up    down
fe-4/1/2.0      up    down
```

Meaning The sample output lists only the Fast Ethernet interfaces. It shows the status of both the physical and logical interfaces. For a description of what the output means, see Table 88 on page 904.

Table 88: Status of Fast Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
fe-2/1/0 Admin Up Link Up	fe-2/1/0.0 Admin Up Link Up	This interface has both the physical and logical links up and running.
fe-3/0/2 Admin Up Link Down	fe-3/0/2.0 Admin Up Link Down	This interface has the physical link down, the link layer down, or both down (**Link Down**). The logical link is also down as a result.
fe-4/1/0 Admin Down Link Up	fe-4/1/0.0 Admin Up Link Down	This interface is administratively disabled and the physical link is healthy (**Link Up**), but the logical interface is not established. The logical interface is down because the physical link is disabled.
fe-4/1/2 Admin Up Link Down	fe-4/1/2.0 Admin Up Link Down	This interface has both the physical and logical links down.

See Also

- Display the Status of Gigabit Ethernet Interfaces on page 904

Display the Status of Gigabit Ethernet Interfaces

Purpose To display the status of Gigabit Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Sample Output Action

```
user@host> show interfaces terse ge*
Interface        Admin Link Proto Local                 Remote
ge-2/2/0         down  down
ge-2/2/0.0       up    down inet  65.113.23.105/30
ge-2/3/0         up    up
ge-2/3/0.0       up    up   inet  65.115.56.57/30
ge-3/1/0         up    up
ge-3/1/0.0       up    up   inet  65.115.56.193/30
ge-3/2/0         up    down
```

Meaning This sample output lists only the Gigabit Ethernet interfaces. It shows the status of both the physical and logical interfaces. See Table 89 on page 905 for a description of what the output means.

Table 89: Status of Gigabit Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
ge-2/2/0 **Admin Down** **Link Down**	**ge-2/2/0.0** **Admin Up** **Link Down**	This interface is administratively disabled (**Admin Down**). Both the physical and logical links are down (**Link Down**).
ge-2/3/0 **Admin Up** **Link Up**	**ge-2/3/0.0** **Admin Up** **Link Up**	This interface has both the physical and logical links up and running.
ge-3/2/0 **Admin Up** **Link Down**	**ge-3/2/0.0** **Admin Up** **Link Down**	This interface has both the physical link and the logical interface down.

See Also

- Display the Status of Fast Ethernet Interfaces on page 903

Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface

Purpose To display the status of a specific Fast Ethernet or Gigabit Ethernet interface when you need to investigate its status further, use the following Junos OS CLI operational mode command:

Action user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**

Sample Output 1

The following sample output is for a Fast Ethernet interface with the physical link up:

```
user@host> show interfaces fe-2/1/0
Physical interface: fe-2/1/0, Enabled, Physical link is Up
  Interface index: 31, SNMP ifIndex: 35
  Description: customer connection
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:86:71:1b, Hardware address: 00:90:69:86:71:1b
  Input rate     : 25768 bps (11 pps), Output rate: 1576 bps (3 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-2/1/0.0 (Index 2) (SNMP ifIndex 43)
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.116.151.218/29, Local: 10.119.115.217
        Broadcast: 10.116.151.225
```

Sample Output 2

The following output is for a Gigabit Ethernet interface with the physical link up:

```
user@host> show interfaces ge-3/1/0
Physical interface: ge-3/1/0, Enabled, Physical link is Up
  Interface index: 41, SNMP ifIndex: 55
  Description: customer connection
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 1000mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:85:71:99, Hardware address: 00:90:69:85:71:99
  Input rate     : 7412216 bps (1614 pps), Output rate: 2431184 bps (1776 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-3/1/0.0 (Index 11) (SNMP ifIndex 57)
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.117.65.192/30, Local: 10.115.65.193
        Broadcast: 10.115.65.195
```

Meaning The first line of sample output 1 and 2 shows that the physical link is up. This means that the physical link is healthy and can pass packets. Further down the sample output, look for active alarms and defects. If you see active alarms or defects, to further diagnose the problem, see Step 3, "Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface" on page 907, to display more extensive information about the Fast Ethernet interface and the physical interface that is down.

Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface

Purpose To display extensive status information about a specific Fast Ethernet or Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**

Sample Output

The following sample output is for a Fast Ethernet interface:

```
user@router> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Up
  Interface index: 47, SNMP ifIndex: 38
  Description: Test
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
  Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
  Traffic statistics:
   Input  bytes  :           373012658                    0 bps
   Output bytes  :           153026154                 1392 bps
   Input  packets:             1362858                    0 pps
   Output packets:             1642918                    3 pps
   Input errors:
    Errors: 0 , Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
    L3 incompletes: 1 , L2 channel errors: 0 , L2 mismatch timeouts: 0
    FIFO errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
    HS link CRC errors: 0, FIFO errors: 0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                     439703575        177452093
    Total packets                      1866532          1642916
    Unicast packets                     972137          1602563
    Broadcast packets                       30             2980
    Multicast packets                   894365            37373
    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                 1866532
    Input packet rejects                     0
     Input DA rejects          503674
     Input SA rejects           0
    Output packet count                                 1642916
```

```
    Output packet pad count                                  0
    Output packet error count                                0
    CAM destination filters: 5, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: OK
    Link partner: Full-duplex, Flow control: None
  PFE configuration:
    Destination slot: 1, Stream number: 15
    CoS transmit queue bandwidth:
      Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
    CoS weighted round-robin:
      Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
    Description: Test
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.115.107.192/29, Local: 10.115.107.193
        Broadcast: 10.115.107.199
```

Meaning The sample output shows where the errors might be occurring and includes autonegotiation information. See Table 90 on page 908 for a description of errors to look for.

Table 90: Errors to Look For

Error	Meaning
Policed discards	Discarded frames that were not recognized or were not of interest.
L2 channel errors	Packets for which the router could not find a valid logical interface. For example, the packet is for a virtual LAN (VLAN) that is not configured on the interface.
MTU	The maximum transmission unit (MTU) must match the interface of either the router at the remote end of the Fast Ethernet or Gigabit Ethernet link, or that of the switch.
Input DA rejects	Number of packets with a destination Media Access Control (MAC) address that is not on the accept list. It is normal to see this number increment.
Input SA rejects	Number of packets with a source MAC address that is not on the accept list. This number only increments when source MAC address filtering is configured.

If the physical link is down, look at the active alarms and defects for the Fast Ethernet or Gigabit Ethernet interface and diagnose the Fast Ethernet or Gigabit Ethernet media accordingly. See “Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters” on page 927 for an explanation of Fast Ethernet and Gigabit Ethernet alarms.

Table 91 on page 909 lists and describes some MAC statistics errors to look for.

Table 91: MAC Statistics Errors

Error	Meaning
CRC/Align errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad FCS with an integral number of octets (**FCS Error**) or a bad FCS with a non-integral number of octets (**Alignment Error**).
MAC control frames	The number of MAC control frames.
MAC pause frames	The number of MAC control frames with **pause** operational code.
Jabber frames	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Note that 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 where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragment frames	The total number of packets received that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error an alignment error. Note that it is entirely normal for fragment frames to increment because both runts (which are normal occurrences due to collisions) and noise hits are counted.

Autonegotiation is the process that connected Ethernet interfaces use to communicate the information necessary to interoperate. Table 92 on page 909 explains the autonegotiation information of the **show interface *interface-name* extensive** command output.

Table 92: Autonegotiation Information

Autonegotiation Field Information	Explanation
Negotiation status: Incomplete	The **Negotiation status** field shows **Incomplete** when the Ethernet interface has the speed or link mode configured.
Negotiation status: No autonegotiation	The **Negotiation status** field shows **No autonegotiation** when the remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
Negotiation status: Complete **Link partner status: OK**	The **Negotiation status** field shows **Complete** and the **Link partner** field shows **OK** when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process completes successfully.
Link partner: Half-duplex	The **Link partner** field can be **Full-duplex** or **Half-duplex** depending on the capability of the attached Ethernet device.
Flow control: Symmetric/asymmetric	The **Flow control** field displays the types of flow control supported by the remote Ethernet device.

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface

Purpose To monitor statistics for a Fast Ethernet or Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action user@host> **monitor interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**

CAUTION: We recommend that you use the **monitor interface fe-*fpc/pic/port*** or **monitor interface ge-*fpc/pic/port*** command only for diagnostic purposes. Do not leave these commands on during normal router operations because real-time monitoring of traffic consumes additional CPU and memory resources.

Sample Output

The following sample output is for a Fast Ethernet interface:

```
user@host> monitor interface fe-2/1/0
Interface: fe-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:                                             Current Delta
  Input bytes:              282556864218 (14208 bps)                  [40815]
  Output bytes:              42320313078 (384 bps)                      [890]
  Input packets:               739373897 (11 pps)                       [145]
  Output packets:              124798688 (1 pps)                         [14]
Error statistics:
  Input errors:                        0                                  [0]
  Input drops:                         0                                  [0]
  Input framing errors:                0                                  [0]
  Policed discards:              6625892                                  [6]
  L3 incompletes:                     75                                  [0]
  L2 channel errors:                   0                                  [0]
  L2 mismatch timeouts:                0                                  [0]
  Carrier transitions:                 1                                  [0]
  Output errors:                       0                                  [0]
  Output drops:                        0                                  [0]
  Aged packets:                        0                                  [0]
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets              464751787                                [154]
  Packet error count                   0                                  [0]
```

Meaning Use the information from this command to help narrow down possible causes of an interface problem.

NOTE: If you are accessing the router from the console connection, make sure you set the CLI terminal type using the **set cli terminal** command.

The statistics in the second column are the cumulative statistics since the last time they were cleared using the **clear interfaces statistics** ***interface-name*** command. The statistics in the third column are the cumulative statistics since the **monitor interface** ***interface-name*** command was executed.

If the input errors are increasing, verify the following:

1. Check the cabling to the router and have the carrier verify the integrity of the line. To verify the integrity of the cabling, make sure that you have the correct cables for the interface port. Make sure you have single-mode fiber cable for a single-mode interface and multimode fiber cable for a multimode interface.
2. For a fiber-optic connection, measure the received light level at the receiver end and make sure that it is within the receiver specification of the Ethernet interface. See "Fiber-Optic Ethernet Interface Specifications" on page 911 for the fiber-optic Ethernet interface specifications.
3. Measure the transmit light level on the Tx port to verify that it is within specification. See "Fiber-Optic Ethernet Interface Specifications" on page 911 for the optical specifications.

Fiber-Optic Ethernet Interface Specifications

Table 93 on page 911 shows the specifications for fiber-optic interfaces for Juniper Networks routers.

Table 93: Fiber-Optic Ethernet Interface Specifications

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
Gigabit Ethernet					
Duplex SC connector					
LH optical interface	49.5-mile 70-km reach on 8.2-micrometer SMF	1480 to 1580 nm	-3 to +2 dBm	-3 dBm	-23 dBm (BER 1012) for SMF
LX optical interface	6.2-mile 10-km reach on 9/125-micrometer SMF 1804.5-ft 550-m reach on 62.5/125- and 50/125-micrometer MMF	1270 to 1355 nm	-11 to -3 dBm	-3 dBm	-19 dBm
SX optical interface	656-ft 200-m reach on 62.5/125-micrometer MMF 1640-ft 500-m reach on 50/125-micrometer MMF	830 to 860 nm	-9.5 to -4 dBm	-3 dBm	-17 dBm
Fast Ethernet 8-Port					

Table 93: Fiber-Optic Ethernet Interface Specifications (continued)

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
FX optical interface with MT-RJ connector	1.24-mile 2-km reach on 62.5/125-micrometer MMF	1270 to 1380 nm	-20 to -14 dBm	-14 dBm	-34 dBm

See Also

- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Understanding Fiber-Optic Cable Signal Loss, Attenuation, and Dispersion*
- *Calculating Power Budget and Power Margin for Fiber-Optic Cables*

Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Purpose To use loopback testing to isolate Fast Ethernet and Gigabit Ethernet interface problems.

Action Table 94 on page 912 provides links and commands for using loopback testing for Fast Ethernet and Gigabit Ethernet interfaces.

Table 94: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
“Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 913	
1. Create a Loopback on page 914	
a. Create a Physical Loopback for a Fiber-Optic Interface on page 914	Connect the transmit port to the receive port.

Table 94: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces (continued)

Tasks	Command or Action
b. Create a Loopback Plug for an RJ-45 Ethernet Interface on page 914	Cross pin 1 (TX+) and pin 3 (RX+) together, and pin 2 (TX-) and pin 6 (RX-) together.
c. Configure a Local Loopback on page 915	**[edit interfaces *interface-name* (fastether-options | gigether-options)]** **set loopback** **show** **commit**
2. Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 916	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
3. Configure a Static Address Resolution Protocol Table Entry on page 919	**show interfaces ge-*fpc/pic/port*** **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address]*** **set arp *ip-address* mac *mac-address* show** **commit** **run show arp no-resolve**
4. Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 923	**clear interfaces statistics fe-*fpc/pic/port* | ge-*fpc/pic/port***
5. Ping the Fast Ethernet or Gigabit Ethernet Interface on page 924	**ping *remote-IP-address* bypass-routing interface (fe-*fpc/pic/port* | ge-*fpc/pic/port* count 100 rapid**
6. Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 925	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
"Diagnose a Suspected Circuit Problem" on page 927	Perform Steps 2 through 8 from "Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 913.

Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface

Problem **Description:** When you suspect a hardware problem, take the following steps to help verify if there is a problem.

Solution To diagnose a suspected hardware problem with the Ethernet interface, follow these steps:

- Create a Loopback on page 914
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 916
- Configure a Static Address Resolution Protocol Table Entry on page 919

- Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 923
- Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 925

Create a Loopback

You can create a physical loopback or configure a local loopback to help diagnose a suspected hardware problem. Creating a physical loopback is recommended because it allows you to test and verify the transmit and receive ports. If a field engineer is not available to create the physical loopback, you can configure a local loopback for the interface. The local loopback creates a loopback internally in the Physical Interface Card (PIC).

1. Create a Physical Loopback for a Fiber-Optic Interface on page 914
2. Create a Loopback Plug for an RJ-45 Ethernet Interface on page 914
3. Configure a Local Loopback on page 915

Create a Physical Loopback for a Fiber-Optic Interface

Action

To create a physical loopback at the port, connect the transmit port to the receive port using a known good fiber cable.

NOTE: **Make sure you use single-mode fiber for a single-mode port and multimode fiber for a multimode port.**

Meaning

When you create and then test a physical loopback, you are testing the transmit and receive ports of the PIC. This action is recommended if a field engineer is available to create the physical loop as it provides a more complete test of the PIC.

See Also

- Create a Loopback Plug for an RJ-45 Ethernet Interface on page 914
- Configure a Local Loopback on page 915
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 916
- Configure a Static Address Resolution Protocol Table Entry on page 919

Create a Loopback Plug for an RJ-45 Ethernet Interface

Action

To create a loopback plug, cross pin 1 (TX+) and pin 3 (RX+) together, and cross pin 2 (TX-) and pin 6 (RX-) together. You need the following equipment to create the loopback:

- A 6-inch long CAT5 cable
- An RJ-45 connector
- A crimping tool

Figure 58 on page 915 illustrates how to create a loopback plug for an RJ-45 Ethernet interface.

Figure 58: RJ-45 Ethernet Loopback Plug

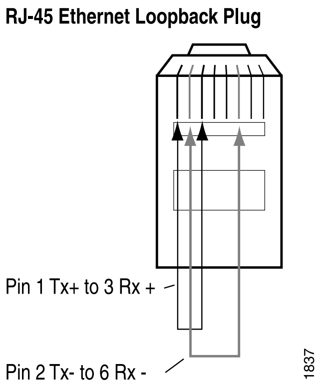

Meaning

When you create and then test a physical loopback, you are testing the RJ-45 interface of the PIC. This action is recommended if a field engineer is available to create the physical loop as it provides a more complete test of the PIC.

See Also

- Configure a Local Loopback on page 915
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 916
- Configure a Static Address Resolution Protocol Table Entry on page 919
-

Configure a Local Loopback

Action

To configure a local loopback without physically connecting the transmit port to the receive port, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   [edit]
   user@host# **edit interfaces** ***interface-name*** (fastether-options | gigether-options)

2. Configure the local loopback:

   [edit interfaces *interface-name* (fastether-options | gigether-options)]
   user@host# **set loopback**

3. Verify the configuration:

   user@host# **show**

   For example:

   ```
   [edit interfaces fe-1/0/0 fastether-options]
   user@host# show
   loopback;
   ```

4. Commit the change:

   user@host# **commit**

   For example:

   ```
   [edit interfaces fe-1/0/0 fastether-options]
   user@host# commit
   commit complete
   ```

When you create a local loopback, you create an internal loop on the interface being tested. A local loopback loops the traffic internally on that PIC. A local loopback tests the interconnection of the PIC but does not test the transmit and receive ports. On an Ethernet interface, you cannot create a remote loopback, therefore there is no option to use a **local** or **remote** statement. Simply including the **loopback** statement at the **[edit interfaces *interface-name* (fastether-options | gigether-options]** hierarchy level, places the interface into local loopback mode.

NOTE: **Remember to delete the loopback statement after completing the test.**

See Also

- Create a Loopback Plug for an RJ-45 Ethernet Interface on page 914
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 916
- Configure a Static Address Resolution Protocol Table Entry on page 919

See Also

- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 916
- Configure a Static Address Resolution Protocol Table Entry on page 919
- Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 923
- Ping the Fast Ethernet or Gigabit Ethernet Interface on page 924
- Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 925

Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up

Purpose Display the status of the Fast Ethernet or Gigabit Ethernet interface to provide the information you need to determine whether the physical link is up or down.

Action To verify that the status of the Fast Ethernet or Gigabit Ethernet interface is up, use the following Junos OS command-line interface (CLI) operational mode command:

user@host> **show interfaces (fe-*fpc/port* | ge-*fpc/pic/port*)**

Sample Output

```
user@host# show interfaces ge-4/0/6 extensive
Physical interface: ge-4/0/6, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 516, Generation: 147
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Loop-Detected
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
  Last flapped   : 2015-01-20 23:40:04 PST (00:02:12 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
  Dropped traffic statistics due to STP State:
   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: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   0                    0                    0

    1 expedited-fo                  0                    0                    0

    2 assured-forw                  0                    0                    0

    3 network-cont                  0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
```



```
    3                    network-control
  Active alarms  : None
  Active defects : None
  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
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
 OK
    Local resolution:
        Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 4
  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
  Interface transmit statistics: Disabled
```

Meaning

The sample output shows that the link is up and there are no alarms in this loopback configuration. When an internal loopback is configured, the physical loopback should come up without an alarm.

Sample Output

When you see that the physical link is down, there may be a problem with the port. The following output is an example of the show interfaces fe-*fpc/pic/port* command when the physical link is down:


```
user@router> show interfaces fe-1/3/0
Physical interface: fe-1/3/0, Enabled, Physical link is Down
  Interface index: 44, SNMP ifIndex: 35
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:db, Hardware address: 00:90:69:8d:2c:db
  Input rate     : 0 bps (0 pps), Output rate: 0 bps (0 pps)
  Active alarms : LINK
  Active defects : LINK
  MAC statistics:
    Input octets: 0, Input packets: 0, Output octets: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Autonegotiation information:
    Negotiation status: Incomplete, Link partner status: Down
    Reason: Link partner autonegotiation failure
    Link partner: Half-duplex, Flow control: None
```

Meaning The sample output shows that the physical link is down and there are active alarms and defects.

Table 95 on page 919 presents problem situations and actions for a physical link that is down.

Table 95: Problems and Solutions for a Physical Link That Is Down

Problem	Action
Cable mismatch	Verify that the fiber connection is correct.
Damaged and/or dirty cable	Verify that the fiber can successfully loop a known good port of the same type.
Too much or too little optical attenuation	Verify that the attenuation is correct per the PIC optical specifications.
The transmit port is not transmitting within the dBm optical range per the specifications	Verify that the Tx power of the optics is within range of the PIC optical specification.
Mismatch between the cable type and the port	Verify that a single-mode fiber cable is connected to a single-mode interface and that a multimode fiber cable is connected to a multimode interface. (This problem does not always cause the physical link to go down; errors and dropped packets are sometimes the result.)

Configure a Static Address Resolution Protocol Table Entry

Purpose

Configure a static Address Resolution Protocol (ARP) entry to allow a packet to be sent out of a looped Ethernet interface.

NOTE: Remove the static ARP entry at the end of the loop test after you have completed the tests and monitored interface traffic.

Action

To configure a static ARP table entry for a Gigabit Ethernet interface, follow these steps. You can follow the same procedure to configure a static ARP entry for a Fast Ethernet interface.

1. Find the Media Access Control (MAC) address for the Gigabit Ethernet interface:

   user@host> **show interfaces ge-*fpc/pic/port***

```
Physical interface: ge-4/0/6, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 516, Generation: 147
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Loop-Detected
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
 Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
  Last flapped   : 2015-01-20 23:40:04 PST (00:13:49 ago)
  Statistics last cleared: 2015-01-20 23:46:15 PST (00:07:38 ago)
  Traffic statistics:
   Input  bytes  :               125500                    0 bps
   Output bytes  :               125482                    0 bps
   Input  packets:                 1281                    0 pps
   Output packets:                 1281                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Dropped traffic statistics due to STP State:
   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: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets     Dropped packets

    0 best-effort                1260                 1260
  0
```



```
    1 expedited-fo                 0                    0
  0
    2 assured-forw                 0                    0
  0
    3 network-cont                 0                    0
  0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                        130624           130624
    Total packets                         1281             1281
    Unicast packets                       1280             1280
    Broadcast packets                        1                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                    1281
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                    1281
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote
fault: OK
    Local resolution:
        Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 4
  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
  Interface transmit statistics: Disabled

  Logical interface ge-4/0/6.0 (Index 72) (SNMP ifIndex 573) (Generation 137)

    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    Traffic statistics:
```

```
        Input  bytes  :              125500
        Output bytes  :              123480
        Input  packets:                1281
        Output packets:                1260
      Local statistics:
        Input  bytes  :                  60
        Output bytes  :                2002
        Input  packets:                   1
        Output packets:                  21
      Transit statistics:
        Input  bytes  :                   0                    0 bps
        Output bytes  :                   0                    0 bps
        Input  packets:                   0                    0 pps
        Output packets:                   0                    0 pps
      Security: Zone: HOST
      Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp
nhrp ospf pgm
      pim rip router-discovery rsvp sap vrrp
      Flow Statistics :
      Flow Input statistics :
        Self packets :                     0
        ICMP packets :                     40
        VPN packets :                      0
        Multicast packets :                0
        Bytes permitted by policy :        107520
        Connections established :          20
      Flow Output statistics:
        Multicast packets :                0
        Bytes permitted by policy :        107520
      Flow error statistics (Packets dropped due to):
        Address spoofing:                  0
        Authentication failed:             0
        Incoming NAT errors:               0
        Invalid zone received packet:      0
        Multiple user authentications:     0
        Multiple incoming NAT:             0
        No parent for a gate:              0
        No one interested in self packets: 0
        No minor session:                  0
        No more sessions:                  0
        No NAT gate:                       0
        No route present:                  11
        No SA for incoming SPI:            0
        No tunnel found:                   0
        No session for a gate:             0
        No zone or NULL zone binding       0
        Policy denied:                     0
        Security association not active:   0
        TCP sequence number out of window: 0
        Syn-attack protection:             0
        User authentication errors:        0
      Protocol inet, MTU: 1500, Generation: 158, Route table: 0
        Flags: Sendbcast-pkt-to-re
        Addresses, Flags: Is-Preferred Is-Primary
          Destination: 10.108.120.0/30, Local: 10.108.120.1, Broadcast:
10.108.120.3,
          Generation: 158
      Protocol multiservice, MTU: Unlimited, Generation: 159, Route table: 0
        Policer: Input: __default_arp_policer__
```

2. In configuration mode, go to the following hierarchy level:

[edit]
user@host# **edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet address** *address*

3. Configure the static ARP entry:

   user@host# **set arp** *ip-address* **mac** *mac-address*

4. Commit the configuration:

   user@host# **commit**

5. Verify that the static ARP entry is installed:

   [edit interfaces ge-4/0/6 unit 0 family inet address 10.108.120.1/30]
   user@host# **run show arp no-resolve**

```
MAC Address       Address         Interface     Flags
00:1f:12:fe:c5:2e 10.108.120.2    ge-4/0/6.0           permanent
52:54:00:7d:33:4c 10.204.128.35   fxp0.0               none
52:54:00:65:11:50 10.204.128.36   fxp0.0               none
52:54:00:da:30:82 10.204.128.37   fxp0.0               none
52:54:00:3a:cf:4b 10.204.128.38   fxp0.0               none
52:54:00:de:88:5f 10.204.128.45   fxp0.0               none
52:54:00:48:03:b7 10.204.128.46   fxp0.0               none
```

Meaning

The sample output is for Step 1 through Step 6 and shows that a static ARP entry was configured on Gigabit Ethernet interface **ge-4/0/6**.

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics

Purpose

You can reset the Fast Ethernet and Gigabit Ethernet interface statistics. Resetting the statistics provides a clean start so that previous input/output errors and packet statistics do not interfere with the current diagnostics.

Action

To clear all statistics for the interface, use the following Junos OS CLI operational mode command:

user@host> **clear interfaces statistics** (**fe-***fpc/pic/port* | **ge-***fpc/pic/port*)

Sample Output

```
user@host> clear interfaces statistics ge-4/0/6
user@host>
```

Meaning

This command clears the interface statistics counters for the Gigabit Ethernet interface only.

Ping the Fast Ethernet or Gigabit Ethernet Interface

Purpose Use the ping command to verify the loopback connection.

Action To send ping packets from the Ethernet interface, use the following Junos OS CLI operational mode command:

user@host> **ping** ***remote-IP-address*** **bypass-routing interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*) count 100 rapid**

Sample Output

```
user@router> ping 10.108.120.2 bypass-routing interface ge-7/2/1 count 100 rapid
PING 10.108.120.2 (10.108.120.2): 56 data bytes
36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e871   0 0000  01  01 cc5c 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e874   0 0000  01  01 cc59 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e878   0 0000  01  01 cc55 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e87c   0 0000  01  01 cc51 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e880   0 0000  01  01 cc4d 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e884   0 0000  01  01 cc49 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
```

Meaning The sample output shows that the time to live (TTL) expired, indicating that the link is receiving the frames from the ping test. The MAC address used is the same as the physical address of the port being tested because this allows the port to accept the frames from the ping test. As the packet is looped over the link, you expect to receive a TLL exceeded message for each ping sent. These messages are generated because the ping packets are repeatedly looped between the router and the physical loopback. When the packet is sent to the other end of the link, which does not exist, the loopback returns the packet back to the same interface, where it is again subjected to the Packet Forwarding Engine fabric for routing. After the route lookup, the TTL is decremented, and the packet is again sent out of the looped interface. This process repeats until the packed is either lost, or the TLL expires with subsequent TTL expired message displayed. Should any errors occur, the packet is discarded and a time-out error is displayed, rather than the expected TTL expired message. Note that the default TTL for ICMP echo packets in Junos OS is 64. This means a given test packet must be successfully sent and received 63 times before a TTL expired message can be generated. You can alter the TTL value to adjust the tolerance for loss, for example, a value of 255 is the most demanding test because now the packet must be sent and received error free 254 times.

Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics

Purpose Persistent interface error statistics indicate that you need to open a case with the Juniper Networks Technical Assistance Center (JTAC).

Action To check the local interface for error statistics, use the following Junos OS CLI operational mode command:

user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**

Sample Output

```
user@router> show interfaces ge-4/0/6 extensive
Physical interface: ge-4/0/6, Enabled, Physical link is Up  Interface index: 144, SNMP
ifIndex: 516, Generation: 147
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Loop-Detected
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
  Last flapped   : 2015-01-20 23:40:04 PST (00:02:12 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
  Dropped traffic statistics due to STP State:
   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: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0
```

```
    2 assured-forw                  0                    0                    0

    3 network-cont                  0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms : None
  Active defects : None
  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
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
       Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
 OK
    Local resolution:
        Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 4
  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
  Interface transmit statistics: Disabled
```

Meaning Check for any error statistics. There should not be any input or output errors. If there are any persistent input or output errors, open a case with the Juniper Networks Technical

Assistance Center (JTAC) at **support@juniper.net**, or at 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).

Diagnose a Suspected Circuit Problem

Purpose When you suspect a circuit problem, it is important to work with the transport-layer engineer to resolve the problem. The transport-layer engineer may create a loop to the router from various points in the network. You can then perform tests to verify the connection from the router to that loopback in the network.

Action After the transport-layer engineer has created the loop to the router from the network, you must verify the connection from the router to the loopback in the network. Follow Step 2 through Step 8 in "Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 913. Keep in mind that any problems encountered in the test indicate a problem with the connection from the router to the loopback in the network.

By performing tests to loopbacks at various points in the network, you can isolate the source of the problem.

Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters

Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

Purpose To locate LINK alarm and major counters associated with Fast Ethernet and Gigabit Ethernet interfaces.

Action Table 96 on page 927 provides links and commands for locating LINK alarm and major counters for Fast Ethernet and Gigabit Ethernet interfaces.

Table 96: Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

Tasks	Command or Action
"Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm" on page 928	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
"Fast Ethernet and Gigabit Ethernet Counters" on page 929	

See Also
- *Ethernet Interfaces Feature Guide for Routing Devices*

Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm

Problem **Description:** To display the Fast Ethernet or Gigabit Ethernet LINK alarm, use the following Junos OS command-line interface (CLI) operational mode command:

Solution user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**

Sample Output

The following sample output is for a Fast Ethernet interface:

```
user@host> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Down
  Interface index: 47, SNMP ifIndex: 38
  Description: Test
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
  Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
  Traffic statistics:
   Input  bytes  :           373012658                    0 bps
   Output bytes  :           153026154                 1392 bps
   Input  packets:             1362858                    0 pps
   Output packets:             1642918                    3 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
    L3 incompletes: 1 , L2 channel errors: 0, L2 mismatch timeouts: 0
    FIFO errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
    HS link CRC errors: 0, FIFO errors: 0
  Active alarms : LINK
  Active defects : LINK
  MAC statistics:                      Receive         Transmit
    Total octets                     439703575        177452093
    Total packets                      1866532          1642916
    Unicast packets                     972137          1602563
    Broadcast packets                       30             2980
    Multicast packets                   894365            37373
    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                 1866532
    Input packet rejects                     0
    Input DA rejects                    503674
    Input SA rejects                         0
    Output packet count                                 1642916
    Output packet pad count                                   0
```



```
    Output packet error count                                 0
    CAM destination filters: 5, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: OK
    Link partner: Full-duplex, Flow control: None
  PFE configuration:
    Destination slot: 1, Stream number: 15
    CoS transmit queue bandwidth:
      Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
    CoS weighted round-robin:
      Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
    Description: Test
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.115.107.192/29, Local: 10.115.107.193
        Broadcast: 10.115.107.199
```

Meaning

The sample output shows where the alarm and other errors might be occurring and any counters that are incrementing. The only alarm associated with Fast Ethernet or Gigabit Ethernet interfaces is the LINK alarm. A LINK alarm indicates a physical problem. To isolate where the physical problem might be occurring, conduct loopback testing. See “Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces” on page 912 for information on conducting a loopback test.

NOTE: Since link status is polled once every second, some items that require fast link down detection, such as Multiprotocol Label Switching (MPLS) fast reroute, take longer to execute.

See Also

- Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters on page 927
- Fast Ethernet and Gigabit Ethernet Counters on page 929

- *Ethernet Interfaces Feature Guide for Routing Devices*

Fast Ethernet and Gigabit Ethernet Counters

Problem **Description:** Table 97 on page 930 shows the major counters that appear in the output for the **show interfaces fe-*fpc/pic/port* extensive** and the **show interfaces ge-*fpc/pic/port* extensive** commands. These counters generally increment when there is a problem with a Fast Ethernet or Gigabit Ethernet interface. In the **Counters** column, the counters are listed in the order in which they are displayed in the output.

Table 97: Major Fast Ethernet and Gigabit Ethernet Counters

Counter	Description	Reason for Increment
Input Errors:		
Errors	The sum of the incoming frame aborts and frame check sequence (FCS) errors.	
Policed discards	The frames discarded by the incoming packet match code.	The frames were discarded because they were not recognized or of interest. Usually, this field reports protocols that the Junos OS does not handle.
Drops	The number of packets dropped by the output queue of the I/O Manager application-specific integrated circuit (ASIC).	If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism.
L3 incompletes	The number of packets discarded due to the packets failing Layer 3 header checks.	This counter increments when the incoming packet fails Layer 3 (usually IPv4) checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment.
L2 channel errors	The errors that occur when the software could not find a valid logical interface (such as **fe-1/2/3.0**) for an incoming frame.	This error increments when, for example, a lookup for a virtual LAN (VLAN) fails.
L2 mismatch timeouts	The count of malformed or short packets.	The malformed or short packets cause the incoming packet handler to discard the frame and be unreadable.
FIFO errors	The number of first in, first out (FIFO) errors in the receive direction as reported by the ASIC on the Physical Interface Card (PIC).	The value in this field should always be 0. If this value is not zero, cabling could be badly organized or the PIC could be broken.
Output Errors		
Errors	The sum of outgoing frame aborts and FCS errors.	
Collisions	The number of Ethernet collisions.	The Fast Ethernet PIC supports only full-duplex operation, so this number should always remain 0. If it is incrementing, there is a software bug.
Drops	The 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.
Aged packets	The number of packets that remained in shared packet SDRAM for so long that the system automatically purged them.	The value in this field should never increment. If it increments, it is probably a software bug or broken hardware.

Table 97: Major Fast Ethernet and Gigabit Ethernet Counters (continued)

Counter	Description	Reason for Increment
HS link FCS errors, FIFO errors	The number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.	The value in this field should always be 0. If it increments, either the FPC or the PIC is broken.
Miscellaneous Counters		
Input DA rejects	The number of packets that the filter rejected because the destination Media Access Control (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 router from the far-end system, either there is a bad Address Resolution Protocol (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 router (which the router is rejecting).
Output packet pad count	The number of packets that 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 Internet Protocol (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	The number of entries in the content-addressable memory (CAM) dedicated to destination and source MAC address filters.	There can be up to 64 source entries. If source filtering is disabled, which is the default, the value for these fields should be 0.

See Also

- Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters on page 927
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Understanding Interfaces on ACX Series Universal Metro Routers*
- *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*

Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State

Problem

Description: 10-Gigabit Ethernet port is stuck in DPC or PIC down state.

Environment: Juniper Networks T Series and MX Series routers. Refer to the related documentation section for more information.

Symptoms: The device has failed to initialize because the Ethernet port is down.

Diagnosis

Try disabling and reenabling the interface and resetting the transceiver and cable. If the interface remains down, it can be stuck in DPC or PIC down state.

Does the router function normally after disabling and reenabling the interface and resetting the transceiver and cable?

Yes:

The system is not stuck in DPC or PIC down state. Disabling and reenabling the interface or resetting the transceiver, and cable resolved the issue.

No:

The interface might be stuck in DPC or PIC down state. Refer to the "To resolve the issue" on page 932 section for recovery options.

1.

Resolution ***To resolve the issue***

From the aforementioned diagnosis, you ascertain that the interface is stuck in DPC or PIC down state.

This is not a hardware defect. Implement one of the following solutions on the backup Routing Engine to resolve this issue:

- Reset the PIC.
- Toggle the framing mode.

1. In configuration mode, go to the **[edit interfaces]** hierarchy level.

   user@host1# **edit interfaces** ***interface name***

2. Toggle the framing mode. In the following configuration, WAN-PHY mode is toggled.

   [edit interfaces *interface-name* is in the et-*fpc/pic/port*
   user@host1# **set framing wan-phy**
   user@host1# **commit**
   user@host1# **framing {**
   user@host1# **wan-phy;**
   user@host1# **}**
   user@host1# **delete framing**
   user@host1# **commit**

3. Reset the PIC (T Series Routers)

   user@host1# **request chassis pic fpc-slot** ***x*** **pic-slot** ***y*** **offline**
   user@host1# **request chassis pic fpc-slot** ***x*** **pic-slot** ***y*** **online**

4. Reset the PIC (MX Series Routers)

   user@host1# **request chassis fpc slot** ***x*** **offline**
   user@host1# **request chassis fpc slot** ***x*** **online**

Related Documentation

- Ethernet Interfaces Overview on page 3
- DPCs Supported on MX240, MX480, and MX960 Routers
- T1600 PICs Supported

PART 5

Configuration Statements and Operational Commands

CHAPTER 37

Configuration Statements (OTN)

alarm (optics-options)

Syntax

```
alarm low-light–alarm {
  (link-down | syslog);
}
```

Hierarchy Level

[edit interfaces *interface-name* optics-options]

Release Information

Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold.

Options

link-down—Drop the 10-Gigabit Ethernet link and marks link as down.

syslog—Write the optics information to the system log.

Required Privilege Level

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

Related Documentation

- Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning on page 432
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475

backward-frr-enable

Syntax (backward-frr-enable | no-backward-frr-enable);

Hierarchy Level [edit interfaces *interface-name* otn-options preemptive-fast-reroute]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Enable or disable backward fast reroute status insertion.

Enable backward fast reroute to insert local pre-forward error correction (FEC) bit error rate (BER) status into transmitted OTN frames, notifying the remote interface. The remote interface can use the information to reroute traffic to a different interface. When you enable backward fast reroute and also enable pre-FEC BER monitoring including the signal-degrade-monitor-enable statement, notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol.

NOTE: **When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the signal-degrade-monitor-enable and backward-frr-enable statements.**

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used. Include the ber-threshold-signal-degrade ***value***, ber-threshold-clear ***value***, and interval ***value*** statements at the **[edit interfaces *interface-name*** otn-options signal-degrade**]** hierarchy level to configure the BER thresholds and time interval. See "Understanding Pre-FEC BER Monitoring and BER Thresholds" on page 477 for more information about pre-FEC BER monitoring and determining BER threshold settings.

Default By default, backward fast reroute insertion is disabled.

Options **backward-frr-enable**—Enable backward fast reroute status insertion.

no-backward-frr-enable—Do not enable backward fast reroute status insertion.

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

Related Documentation

- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

ber-threshold-clear

Syntax ber-threshold-clear *value*;

Hierarchy Level [edit interfaces *interface-name* otn-options signal-degrade]
[edit interfaces *interface-name* otn-options odu-signal-degrade]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Specify bit error rate (BER) threshold to clear the interface alarm for signal degradation.

You can configure the BER clear threshold to customize the BER that will clear an interface alarm when signal degrade monitoring is enabled.

NOTE: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See "Understanding Pre-FEC BER Monitoring and BER Thresholds" on page 477 for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

NOTE: You can configure ODU BER thresholds only at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level on the P2-100GE-OTN PIC.

Table 98 on page 941 shows the default values for pre-FEC BER and ODU BER signal degrade threshold values for different PICs. If the BER signal degrade threshold is not configured, the default value is used.

Table 98: Default Clear Threshold Values

PIC or MPC	Default Pre-FEC BER Clear Threshold Value	Default ODU BER Clear Threshold Value
P1-PTX-2-100G-WDM	3.0E–3	Not supported
P2-100GE-OTN	3.0E–3	1.0E–9
P1-PTX-24-10G-W-SFPP	3.0E–3	Not supported
MIC6-100G-CFP2	1.0E-6	1.0E-9
MPC5E	1.0E-6	1.0E-9

To configure the threshold that raises the signal degrade alarm, include the ber-threshold-signal-degrade statement at the same hierarchy level. To configure the time interval during which the BER must stay above or below the configured thresholds to raise or clear the alarm, include the interval statement at the same hierarchy level.

NOTE: For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

Options

Values: ***value***—BER threshold for clearing the signal degradation in scientific notation. Both the mantissa and exponent are configurable. Enter the value in the format *x*E-*n*, where *x* is the mantissa and *n* is the exponent. For example, 4.5E-3.

Range: The mantissa must be a decimal number. There is no limit on the number of digits before or after the decimal point.The exponent must be an integer from 0 through 9.

Default: See Table 98 on page 941 for the default values.

BEST PRACTICE: Always set the **ber-threshold-clear** ***value*** lower than the **ber-threshold-signal-degrade** *value*. For the FEC limits, see the table describing the signal degrade and clear thresholds after configuration in "Understanding Pre-FEC BER Monitoring and BER Thresholds" on page 477.

NOTE: In Junos OS Release 13.2R1, only the exponent is valid input for the BER threshold value, and the mantissa is not configurable. The BER threshold value is 1.0E–*n* where *n* > 0, and the valid range of *n* is from 1 through 10.

Required Privilege Level

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

Related Documentation

- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

ber-threshold-signal-degrade

Syntax ber-threshold-signal-degrade *value*;

Hierarchy Level [edit interfaces *interface-name* otn-options signal-degrade]
[edit interfaces *interface-name* otn-options odu-signal-degrade]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Specify the bit error rate (BER) threshold to raise an interface alarm for signal degradation.

You can configure the BER signal degrade threshold to customize the BER that will raise an interface alarm when signal degrade monitoring is enabled.

NOTE: Configuring a high BER threshold for signal degradation and a long interval might cause the internal bit error counter register to get saturated. For example, for the P1-PTX-2-100G-WDM PIC, the internal bit error counter gets saturated when the error count reaches 2E+29. Therefore, the value of ber-threshold-signal-degrade * line rate / interval must be less than 2E+29 to avoid saturation. Assuming a fixed PIC line rate of 1.27E+11 bits per second and an interval of 1000 ms, the **ber-threshold-signal-degrade** value must be less than 4.22E–3.

If the value of the ber-threshold-signal-degrade * line rate / interval exceeds the saturation limit, the configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See “Understanding Pre-FEC BER Monitoring and BER Thresholds” on page 477 for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

NOTE: You can configure ODU BER thresholds only at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level on the P2-100GE-OTN PIC.

Table 99 on page 944 shows the default values for pre-FEC BER and ODU BER signal degrade threshold values for different PICs. If the BER signal degrade threshold is not configured, the default value is used.

Table 99: Default Signal Degrade Threshold Values

PIC or MPC	Default Pre-FEC BER Signal Degrade Threshold Value	Default ODU BER Signal Degrade Threshold Value
P1-PTX-2-100G-WDM	7.5E–3	Not supported
P2-100GE-OTN	7.5E–3	1.0E–6
P1-PTX-24-10G-W-SFPP	7.5E–3	Not supported
MIC6-100G-CFP2	1.14E-5	1.0E-06
MPC5E	1.14E-5	1.0E-06

To configure the threshold that clears the signal degrade alarm, include the **ber-threshold-clear** statement at the same hierarchy level. To configure the time interval during which the BER must stay above or below the configured thresholds to raise or clear the alarm, include the **interval** statement at the same hierarchy level.

NOTE: For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

Options

value—BER threshold for signal degradation in scientific notation. Both the mantissa and exponent are configurable. Enter the value in the format *x*E–*n*, where *x* is the mantissa and *n* is the exponent. For example, 4.5E–3.

Range: The mantissa must be a decimal number. There is no limit on the number of digits before or after the decimal point.The exponent must be an integer from 0 through 9.

Default: See Table 99 on page 944.

NOTE: In Junos OS Release 13.2R1, only the exponent is valid input for the BER threshold value, the mantissa is not configurable. The BER threshold value is 1.0E–*n* where *n* > 0, and the valid range of *n* is from 1 through 10.

BEST PRACTICE: To enable proactive protection before packet loss occurs, set the **ber-threshold-signal-degrade** *value* below the FEC limit. For the FEC limits, see the table describing the signal degrade and clear thresholds after configuration in "Understanding Pre-FEC BER Monitoring and BER Thresholds" on page 477.

Required Privilege Level

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

Related Documentation

- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

bypass

Syntax (bypass | no-bypass);

Hierarchy Level [edit interfaces *interface-name* otn-options odu-delay-management]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Pass or do not pass the delay measurement (DM) value through a node.

Default If you omit the bypass statement, the default behavior is to disable ODU delay management options.

By default, do not pass the DM value through a node.

Options **bypass**—Pass the DM value through a node.

no-bypass—Do not pass the DM value through a node.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

bytes (otn-options)

Syntax bytes transmit-payload-type *value*;

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify the transmit payload type on OTN header bytes.

Options *value*—Transmit payload type.

Range: 0 through 255 bytes

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

fec

Syntax fec (efec | gfec | gfec-sdfec |hgfec | sd-fec | ufec | none);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement and **gfec-sdfec** option introduced in Junos OS Release 13.2 for PTX Series routers. with P1-PTX-2-100G-WDM PIC.
Options **efec**, **gfec**, and **ufec** introduced in Junos OS Release 13.3 for MX Series routers. with MPC5E-100G10G, MPC5E-40G10G, MIC6-10G-OTN, and MIC6-100G-CFP2.
Options **efec**, **gfec**, and **ufec** introduced in Junos OS Release 14.1 for PTX Series routers. with P1-PTX-24-10G-W-SFPP.
Option **hgfec** introduced in Junos OS Release 15.1F5 for MX Series Routers with MIC3-100G-DWDM MIC.
Option **sdfec** introduced in Junos OS Release 15.1F5 for MX Series Routers with MIC3-100G-DWDM MIC.
Option **sdfec** introduced in Junos OS Release 15.1F6 for PTX Series Routers with PTX-5-100G-WDM PIC.

Description Enable forward error correction (FEC) mode.

Default If you do not specify a mode, the default mode is **gfec**. On PTX Series routers with P1-PTX-2-100G-WDM, the default value is **gfec-sdfec**. On PTX Series routers with PTX-5-100G-WDM and on MX Series routers with MIC3-100G-DWDM, the default value is **sdfec**.

Options efec—(M Series, MX Series routers and PTX Series routers only) G.975.1 I.4 enhanced forward error correction (EFEC) is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

gfec—(M series, MX Series routers and PTX Series routers only) G.709 generic forward error correction (GFEC) mode is configured to detect and correct bit errors.

gfec-sdfec—(PTX Series routers only) GFEC and soft-decision forward error correction (SD-FEC) modes are configured to detect and correct bit errors.

hgfec—(MX Series routers only) High gain forward error correction mode is configured to detect and correct bit errors.

sdfec—(MX Series routers and PTX Series routers only) Soft-decision forward error correction mode is configured to detect and correct bit errors.

none—(M Series and MX Series routers only) FEC mode is not configured.

NOTE: On MX Series routers with MIC3-100G-DWDM and PTX Series routers with PTX-5-100G-WDM, none option is not supported. The fec mode must be enabled on the MIC3-100G-DWDM MIC and the PTX-5-100G-WDM PIC.

ufec—(MX Series routers and PTX Series routers only) G.975.1 I.7 Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

Required Privilege Level

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482
- Supported Forward Error Correction Modes on MX Series Routers on page 481
- Supported Forward Error Correction Modes on PTX Series Routers on page 481
- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477

fixed-stuff-bytes

Syntax (fixed-stuff-bytes | no-fixed-stuff-bytes);

Hierarchy Level [edit interfaces *interface-name* otn-options rate]

Release Information Statement introduced in Junos OS Release 9.4.

Description Enable or disable fixed stuff bytes.

Default By default, no fixed stuff bytes are set.

Options **fixed-stuff-bytes**—Fixed stuff bytes 11.0957 Gbps.

no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps.

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

interval

Syntax interval *value*;

Hierarchy Level [edit interfaces *interface-name* otn-options signal-degrade]
[edit interfaces *interface-name* otn-options odu-signal-degrade]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Specify the interval for which the BER must stay above the signal degradation threshold—as configured in the ber-threshold-signal-degrade ***value*** statement—for the alarm to raised. After an alarm is raised, if the BER returns below the clear threshold—as configured in the ber-threshold-clear ***value*** statement—for the specified interval, the alarm is cleared.

NOTE: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See “Understanding Pre-FEC BER Monitoring and BER Thresholds” on page 477 for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

NOTE: You can configure ODU BER thresholds only at the [edit interfaces *interface-name* otn-options odu-signal-degrade] hierarchy level on the P2-100GE-OTN PIC.

Options ***value***—Time interval in milliseconds.

NOTE: For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms * 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

NOTE: For P1-PTX-24-10G-W-SFPP PIC and P2-100GE-OTN PIC, when the router cannot configure BER with the given interval, it selects an optimum interval that is supported for the given BER configuration. If the router is still not able to support the configuration (for example, with a wider gap between the degrade set and clear values), the default values are used and a log is generated.

For the P2-10G-40G-QSFPP PIC, the time interval is supported in multiples of 100 ms. For example, when you configure the interval as 10 ms, then it is rounded off to the nearest multiple of 100 ms.

Range: 1 ms through 1000 ms.

Default: 100 ms.

NOTE: For the P2-100GE-OTN PIC, the default value is 10 ms.

Required Privilege Level

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

is-ma

Syntax (is-ma | no-is-ma);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether masked alarms are enabled or disabled.

Default If you omit the is-ma statement, masked alarms are disabled.

Options **is-ma**—Enable masked alarms.

no-is-ma—Do not enable masked alarms.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

laser-enable

Syntax (laser-enable | no-laser-enable);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether lasers are enabled or disabled.

Default If you omit the laser-enable statement, lasers are disabled.

Options **laser-enable**—Enable lasers.

no-laser-enable—Do not enable lasers.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

line-loopback

Syntax (line-loopback-enable | no-line-loopback);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether line-loopback is enabled or disabled.

Default If you omit the line-loopbback-enable statement, line-loopback is disabled.

Options **line-loopback-enable**—Enable line-loopback.

no-line-loopback—Disable line-loopback.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

local-loopback

Syntax (local-loopback-enable | no-local-loopback);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether local-loopback is enabled or disabled.

Default If you omit the local-loopbback-enable statement, local-loopback is disabled.

Options **local-loopback-enable**—Enable local-loopback.

no-local-loopback—Disable local-loopback.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

monitor-end-point

Syntax (monitor-end-point | no-monitor-end-point);

Hierarchy Level [edit interfaces *interface-name* otn-options odu-delay-management]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Originate or do not originate the connection monitor end point.

Default By default, do not originate the connection monitor end point.

Options **monitor-end-point**—Originate the connection monitor end point.

no-monitor-end-point—Do not originate the connection monitor end point.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

no-odu-backward-frr-enable

Syntax no-odu-backward-frr-enable;

Hierarchy Level [edit interfaces *interface-name* otn-options preemptive-fast-reroute]

Release Information Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Disable preemptive fast reroute (FRR) ODU backward FRR insertion.

Default By default, FRR ODU backward FRR insertion is disabled.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

no-odu-signal-degrade-monitor-enable

Syntax no-odu-signal-degrade-monitor-enable;

Hierarchy Level [edit interfaces *interface-name* otn-options preemptive-fast-reroute]

Release Information Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Disable monitoring of signal degradation of ODU BER in the received OTN frames.

Default By default, FRR signal degrade monitoring disabled.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

number-of-frames

Syntax number-of-frames *value*;

Hierarchy Level [edit interfaces *interface-name* otn-options odu-delay-management]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify the number of consequent frames to declare a delay measurement (DM) session completed.

Options *value*—Number of consequent frames to declare DM completed.

Range: 0 through 255 frames.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

oc192

Syntax oc192;

Hierarchy Level [edit interfaces *interface-name* otn-options rate]

Release Information Statement introduced in Junos OS Release 13.3 for MX Series routers.

Description Set the line rate or speed of the OTN signal to optical channel transport unit 2 (OTU2).

Options **oc192**—OTU2 line rate or 10 Gbps

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- Configuring 100-Gigabit DWDM OTN PICs on page 482

odu-delay-management

Syntax

```
odu-delay-management {
  (bypass | no-bypass);
  (monitor-end-point | no-monitor-end-point);
  number-of-frames value;
  (no-start-measurement | start-measurement;
}
```

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify Optical Channel Data Unit (ODU) delay management options.

Default If you omit the odu-delay-management statement, the ODU delay management options are disabled.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

odu-backward-frr-enable

Syntax odu-backward-frr-enable;

Hierarchy Level [edit interfaces *interface-name* otn-options preemptive-fast-reroute]

Release Information Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

Default By default, FRR ODU backward FRR insertion is disabled.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

odu-signal-degrade

Syntax

```
odu-signal-degrade {
  ber-threshold-clear;
  ber-threshold-signal-degrade;
  interval
}
```

Hierarchy Level

[edit interfaces *interface-name* otn-options]

Release Information

Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description

Specify optical channel data unit (ODU) signal degradation threshold-related values.

Default

If you omit the **odu-signal-degrade** statement, the default threshold values are used.

The following are the default threshold values for optical channel data unit (ODU) signal degradation for the P2-100GE-OTN PIC:

- ber-threshold-clear—1E-09
- ber-threshold-signal-degrade—1E-06
- interval—10 ms

The following are the default threshold values for optical channel data unit (ODU) signal degradation for the MPC5E and the MIC6-100G-CFP2 MIC:

- ber-threshold-clear—1.14E-5
- ber-threshold-signal-degrade—1.0E-6
- interval—10 ms

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

odu-signal-degrade-monitor-enable

Syntax odu-signal-degrade-monitor-enable;

Hierarchy Level [edit interfaces *interface-name* otn-options preemptive-fast-reroute]

Release Information Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Enable monitoring of signal degradation of ODU BER in the received OTN frames.

Default By default, FRR signal degrade monitoring disabled.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

odu-ttim-action-enable

Syntax (odu-ttim-action-enable | no-odu-ttim-action-enable);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether consequent action for Optical Channel Data Unit (ODU) TTIM is enabled or disabled.

Default If you omit the odu-ttim-action-enable statement, consequent action for ODU TTIM is disabled.

Options **odu-ttim-action-enable**—Enable consequent action for ODU TTIM.

no-odu-ttim-action-enable—Disable consequent action for ODU TTIM.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

otu-ttim-action-enable

Syntax (otu-ttim-action-enable | no-otu-ttim-action-enable);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether consequent action for Optical Channel Transport Unit (OTU) TTIM is enabled or disabled.

Default If you omit the otu-ttim-action-enable statement, consequent action for OTU TTIM is disabled.

Options **otu-ttim-action-enable**—Enable consequent action for OTU TTIM.

no-otu-ttim-action-enable—Disable consequent action for OTU TTIM.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

otu4

Syntax otu4;

Hierarchy Level [edit interfaces *interface-name* otn-options rate]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 13.3 for MX Series routers.

Description Sets the line rate or speed of the OTN signal to optical channel transport unit 4 (OTU4).

Default By default, the rate is OTU4 on PTX Series routers.

Options **otu4**—OTU4 line rate or 100 Gbps

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

pass-through

Syntax (pass-through | no-pass-through);

Hierarchy Level [edit interfaces *interface-name* otn-options rate]

Release Information Statement introduced in Junos OS Release 9.4.

Description Enable or disable OTN pass-through mode.

Default By default, OTN pass-through mode is disabled.

Options **no-pass-through**—Do not enable OTN pass-through mode.

pass-through—Enable OTN pass-through mode.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

prbs

Syntax (prbs | no-prbs);

Hierarchy Level [edit interfaces *interface-name* **otn-options**]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify whether OTN payload Pseudo-Random Binary Sequence (PBRS) is enabled or disabled.

Default By default, OTN payload prbs is disabled.

Options **prbs**—Enable OTN payload PBRS.

no-prbs—Disable OTN payload PBRS.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

preemptive-fast-reroute

Syntax

```
preemptive-fast-reroute {
  (backward-frr-enable | no-backward-frr-enable);
  (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
  (odu-backward-frr-enable | no-odu-backward-frr-enable);
  (odu-signal-degrade-monitor-enable |no-odu-signal-degrade-monitor-enable);
}
```

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 13.3 for MX Series routers.

Description Enable or disable preemptive fast reroute options.

Default By default, backward fast reroute insertion and signal degradation monitoring are disabled.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

rate

Syntax

```
rate {
  (fixed-stuff-bytes | no-fixed-stuff-bytes);
  otu4; oc192;
  (pass-through | no-pass-through);
}
```

Hierarchy Level

[edit interfaces *interface-name* otn-options]

Release Information

Statement introduced in Junos OS Release 9.4.
Statement and **otu4** option introduced in Junos OS Release 13.2 for PTX Series routers.
Option **oc192** introduced in Junos OS Release 13.3 for MX Series routers.

Description

Specify the line rate or speed of the OTN signals.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

remote-loop-enable

Syntax (remote-loop-enable | no-remote-loop-enable);

Hierarchy Level [edit interfaces *interface-name* otn-options odu-delay-management]

Release Information Statement introduced in Junos OS Release 17.1 for 100-Gigabit OTN DWDM PIC with CFP2 on PTX3000 and PTX5000 routers.
Statement introduced in Junos OS Release 17.1 for 100-Gigabit OTN DWDM MIC with CFP2-ACO on MX240, MX480, MX960, MX2010, and MX2020 routers with MPC3E and MPC3E-NG.

Description Enable the remote interface to loop back the delay measurement pattern to the local interface. Delay is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (**DM**) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end (local interface).

NOTE: Do not enable remote loopback on both ends (local and remote). If you enable remote loopback on both interfaces, the delay measurement pattern is looped back continuously between the two interfaces.

Default Delay measurement is disabled by default.

Options **remote-loop-enable**—Enables loopback of the delay measurement pattern at the remote interface.

no-remote-loop-enable—Disables loopback of the delay measurement pattern at the remote interface.

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- *100-Gigabit DWDM OTN MIC with CFP2-ACO*
- *100-Gigabit DWDM OTN PIC with CFP2-ACO (PTX Series)*
- Configuring OTN Interfaces on MIC3-100G-DWDM MIC on page 511
- Configuring OTN Interfaces on PTX-5-100G-WDM PIC on page 519
- Disabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 527

- Enabling ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 525
- Understanding ODU Path Delay Measurement on OTN Networks for Performance Monitoring on page 524
- Understanding the MIC3-100G-DWDM MIC on page 508
- Understanding the PTX-5-100G-WDM PIC on page 516

signal-degrade

Syntax

```
signal-degrade {
  ber-threshold-clear value;
  ber-threshold-signal-degrade value;
  interval value;
}
```

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 13.3 for MX Series routers.
.

Description Specify bit error rate (BER) signal degradation thresholds and time interval for raising and clearing alarms for optical transport network (OTN) links.

Default If you omit the **signal-degrade** statement, the default threshold values are used.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

signal-degrade-monitor-enable

Syntax (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);

Hierarchy Level [edit interfaces *interface-name* otn-options preemptive-fast-reroute]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Enable or disable pre-forward error correction (FEC) bit error rate (BER) monitoring.

With pre-FEC BER monitoring enabled, when the configured pre-FEC BER signal degrade threshold is reached, the PIC stops forwarding packets to the remote interface and raises an interface alarm. Ingress packets continue to be processed. If pre-FEC BER monitoring is used with MPLS fast reroute or another link protection method, then traffic is rerouted to a different interface.

You can also configure backward fast reroute to insert local pre-FEC BER status into transmitted OTN frames, notifying the remote interface of signal degradation. The remote interface can use the information to reroute traffic to a different interface. If you use pre-FEC BER monitoring together with backward fast reroute, then notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol. To configure backward fast reroute, include the backward-frr-enable statement at the same hierarchy level.

NOTE: **When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the signal-degrade-monitor-enable and backward-frr-enable statements.**

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used. Include the ber-threshold-signal-degrade ***value***, ber-threshold-clear ***value***, and interval ***value*** statements at the **[edit interfaces *interface-name*** otn-options signal-degrade**]** hierarchy level to configure the BER thresholds and time interval. See “Understanding Pre-FEC BER Monitoring and BER Thresholds” on page 477 for more information about pre-FEC BER monitoring and determining BER threshold settings.

Default By default, pre-FEC BER signal degrade monitoring is disabled.

Options **signal-degrade-monitor-enable**—Enable pre-FEC BER signal degrade monitoring.

no-signal-degrade-monitor-enable—Do not enable pre-FEC BER signal degrade monitoring.

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

Related Documentation
- Understanding Pre-FEC BER Monitoring and BER Thresholds on page 477
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

start-measurement

Syntax (no-start-measurement | start-measurement);

Hierarchy Level [edit interfaces *interface-name* otn-options odu-delay-management]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Start or do not start a delay measurement (DM) session.

Default By default, do not start a DM session.

Options **no-start-measurement**—Do not start a DM session.

start-measurement—Start a DM session.

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

Related Documentation
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

tca

Syntax tca *tca-identifier* (enable-tca | no-enable-tca) (threshold *number* | threshold-24hrs *number*)

Hierarchy Level [edit interfaces *interface-name* optics-options]

[edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 14.2 on the PTX Series.

Description TCAs can give the management system an early indication as to the state of the associated entity when it crosses a certain threshold. TCAs can be set for both minimum and maximum values for gauges and only maximum values for counters. The timely detection of TCAs is essential to proactively manage the interface. TCAs are not an indication of a fault, but rather an indication that the entity may be close to a fault. You can choose which TCAs you want monitored by enabling the TCA. You can either keep the default threshold settings or change the settings.

Enable threshold crossing alerts (TCAs) for the following:

- Optical channel data unit (ODU)
- Optical channel transport unit (OTU)
- Laser power
- Laser temperature

Default By default, TCAs are not enabled.

Options *tca-identifier* —At the otn-options hierarchy level, it can be one of the following:

- **odu-tca-bbe**—ODU background block error threshold-crossing defect trigger
- **odu-tca-bbe-fe**—ODU far-end background block error threshold-crossing defect trigger
- **odu-tca-es**—ODU errored seconds threshold-crossing defect trigger
- **odu-tca-es-fe**—ODU far-end errored seconds threshold-crossing defect trigger
- **odu-tca-ses**—ODU severely errored seconds threshold-crossing defect trigger
- **odu-tca-ses-fe**—ODU far-end severely errored seconds threshold-crossing defect trigger
- **odu-tca-uas**—ODU unavailable seconds threshold-crossing defect trigger
- **odu-tca-uas-fe**—ODU far-end unavailable seconds threshold-crossing defect trigger
- **otu-tca-bbe**—OTU background block error threshold-crossing defect trigger
- **otu-tca-bbe-fe**—OTU far-end background block error threshold-crossing defect trigger
- **otu-tca-es**—OTU errored seconds threshold-crossing defect trigger

- **otu-tca-es-fe**—OTU far-end errored seconds threshold-crossing defect trigger
- **otu-tca-fec-ber**—OTU forward error correction bit error rate threshold-crossing defect trigger
- **otu-tca-ses**—OTU severely errored seconds threshold-crossing defect trigger
- **otu-tca-ses-fe**—OTU far-end severely errored seconds threshold-crossing defect trigger
- **otu-tca-uas**—OTU unavailable seconds threshold-crossing defect trigger
- **otu-tca-uas-fe**—OTU far-end unavailable seconds threshold-crossing defect trigger

tca-identifier —At the optics-options hierarchy level, it can be one of the following:

- **carrier-frequency-offset-high-tca**—Carrier frequency high threshold setting trigger
- **carrier-frequency-offset-low-tca**—Carrier frequency low threshold setting trigger
- **rx-power-high-tca**—Rx power high threshold setting trigger
- **rx-power-low-tca**—Rx power low threshold setting trigger
- **temperature-high-tca**—Temperature high threshold setting trigger
- **temperature-low-tca**—Temperature low threshold setting trigger
- **tx-power-high-tca**—Tx power high threshold setting trigger
- **tx-power-low-tca**—Tx power low threshold setting trigger

enable-tca | no-enable-tca—To enable or disable the threshold crossing alert.

threshold | threshold-24hrs:

- threshold *number*—Set the 15-minute interval threshold.
- threshold-24hrs *number*—Set the 24-hour interval threshold.

Required Privilege Level

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

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482
- **optics-options on page 1288**

transport-monitoring

Syntax transport-monitoring;

Hierarchy Level [edit interfaces]

Release Information Statement introduced in Junos OS Release 14.2 for PTX5000 and PTX3000 routers.

Description Monitor the performance and state of packet transport for OTN and optics modules. The following statistics are monitored:

- Packet transport for ninety-six 15–minute intervals for the current 24 hours.
- Cumulative data of the current 24 hours.
- Cumulative data of the previous 24 hours.

If this statement is configured, transport monitoring related information is shown in the output of **show interface transport** command and corresponding MIBs are available. If this option is disabled, an error is shown in the output and corresponding MIBs are not available.

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

Related Documentation

- show interfaces transport pm on page 1994
- *Managed Objects for the OTN Interface*
- *Managed Objects for the Optics Interface*
- *Optical Parameters Associated with Black Link End Points*

trigger

Syntax trigger *trigger-identifier* (hold-time *hold-time-value* | ignore);

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specify defect triggers.

Default By default, triggers are ignored.

Options *trigger-identifier*—(For M Series, MX Series, SRX Series, and T Series routers only) Trigger identifier. It can be one of the following:

- **oc-lof**—Optical channel Loss of Frame defect trigger.
- **oc-lom**—Optical channel Loss of Multiframe defect trigger.
- **oc-los**—Optical channel Loss of Signal defect trigger.
- **oc-wavelength-lock**—Optical channel Wavelength Lock defect trigger.
- **odu-ais**—Optical channel data unit (ODU) Alarm Indication Signal defect trigger.
- **odu-bbe-th**—ODU Background Block Error Threshold defect trigger.
- **odu-bdi**—ODU Backward Defect Indication defect trigger.
- **odu-bei**—(MX Series routers only) ODU Backward Error Indication defect trigger.
- **odu-es-th**—ODU Errored Seconds Threshold defect trigger.
- **odu-iae**—(MX Series routers only) ODU Incoming Alignment Error defect trigger.
- **odu-lck**—ODU Locked defect trigger.
- **odu-oci**—ODU Open Connection Indication defect trigger.
- **odu-sd**—ODU Signal Degrade defect trigger.
- **odu-ses-th**—ODU Severely Errored Seconds Threshold defect trigger.
- **odu-tca-es**—(MX Series routers only) ODU Errored Seconds Threshold crossing defect trigger.
- **odu-tca-ses**—(MX Series routers only) ODU Severely Errored Seconds Threshold crossing defect trigger.
- **odu-tca-uas**—(MX Series routers only) ODU Unavailable Seconds Threshold crossing defect trigger.
- **odu-ttim**—ODU Trail Trace Identifier Mismatch defect trigger.
- **opu-ptim**—(MX Series routers only) Payload Type Identifier Mismatch defect trigger.

- **odu-uas-th**—ODU Unavailable Seconds Threshold defect trigger.
- **opu-ptm**—Optical Channel Payload (OPU) Payload Type Mismatch defect trigger.
- **otu-ais**—Optical Channel Transport Unit (OTU) Alarm Indication Signal defect trigger.
- **otu-bbe-th**—OTU Background Block Error Threshold defect trigger.
- **otu-bdi**—OTU Backward Defect Indication defect trigger.
- **otu-es-th**—OTU Errored Seconds Threshold defect trigger.
- **otu-fec-deg**—OTU FEC Degrade defect trigger.
- **otu-fec-exe**—OTU FEC Excessive Error defect trigger.
- **otu-iae**—OTU Incoming Alignment defect trigger.
- **otu-sd**—OTU Signal Degrade defect trigger.
- **otu-ses-th**—OTU Severely Errored Seconds Threshold defect trigger.
- **otu-tca-es**—(MX Series routers only) OTU Errored Seconds Threshold crossing defect trigger.
- **otu-tca-ses**—(MX Series routers only) OTU Severely Errored Seconds Threshold crossing defect trigger.
- **otu-tca-uas**—(MX Series routers only) OTU Unavailable Seconds Threshold crossing defect trigger.
- **otu-ttim**—OTU Trail Trace Identifier Mismatch defect trigger.
- **otu-uas-th**—OTU Unavailable Seconds Threshold defect trigger.

trigger-identifier—(For PTX Series routers only) Trigger identifier. It can be one of the following:

- **oc-lof**—Optical channel Loss of Frame defect trigger.
- **oc-lom**—Optical channel Loss of Multiframe defect trigger.
- **oc-los**—Optical channel Loss of Signal defect trigger.
- **oc-tsf**—Optical channel TOE security functionality (TSF) defect trigger.
- **oc-wavelength-lock**—Optical channel Wavelength Lock defect trigger.
- **odu-ais**—ODU Alarm Indication Signal defect trigger.
- **odu-bdi**—ODU Backward Defect Indication defect trigger.
- **odu-bei**—ODU Backward Error Indication defect trigger.
- **odu-iae**—ODU IAE defect trigger.
- **odu-lck**—ODU Locked defect trigger.
- **odu-oci**—ODU Open Connection Indication defect trigger.
- **odu-sd**—ODU Signal Degrade defect trigger.
- **odu-tca-bbe**—ODU Background Block Error Threshold crossing defect trigger.
- **odu-tca-bbe-fe**—ODU far-end Background Block Error (BEI) Threshold crossing defect trigger.
- **odu-tca-es**—ODU Errored Seconds Threshold crossing defect trigger.
- **odu-tca-es-fe**—ODU far-end Errored Seconds Threshold crossing defect trigger.
- **odu-tca-ses**—ODU Severely Errored Seconds Threshold crossing defect trigger.
- **odu-tca-ses-fe**—ODU far-end Severely Errored Seconds Threshold crossing defect trigger.
- **odu-tca-uas**—ODU Unavailable Seconds Threshold crossing defect trigger.
- **odu-tca-uas-fe**—ODU far-end Unavailable Seconds Threshold crossing defect trigger.
- **odu-ttim**—ODU Trail Trace Identifier Mismatch defect trigger.
- **opu-ptim**—Payload Type Identifier Mismatch defect trigger.
- **otu-ais**—OTU Alarm Indication Signal defect trigger.
- **otu-bdi**—OTU Backward Defect Indication defect trigger.
- **otu-fec-deg**—OTU FEC Degrade defect trigger.
- **otu-fec-exe**—OTU FEC Excessive Error defect trigger.
- **otu-iae**—OTU Incoming Alignment defect trigger.
- **otu-sd**—OTU Signal Degrade defect trigger.
- **otu-tca-bbe**—OTU Background Block Error Threshold crossing defect trigger.
- **otu-tca-bbe-fe**—OTU far-end Background Block Error (BEI) Threshold crossing defect trigger.
- **otu-tca-es**—OTU Errored Seconds Threshold crossing defect trigger.

- **otu-tca-es-fe**—OTU far-end Errored Seconds Threshold crossing defect trigger.
- **otu-tca-ses**—OTU Severely Errored Seconds Threshold crossing defect trigger.
- **otu-tca-ses-fe**—OTU far-end Severely Errored Seconds Threshold crossing defect trigger.
- **otu-tca-uas**—OTU Unavailable Seconds Threshold crossing defect trigger.
- **otu-tca-uas-fe**—OTU far-end Unavailable Seconds Threshold crossing defect trigger.
- **otu-ttim**—OTU Trail Trace Identifier Mismatch defect trigger.

hold-time *hold-time-value*—Hold time value. It can be one of the following:

- **down**—Delay before marking interface down when defect occurs (1..65534 milliseconds).
- **up**—Delay before marking interface up when defect is absent (1..65534 milliseconds).

NOTE: The trigger hold time value alone does not mark an interface to be up when the defect is absent or mark an interface to be down when the defect occurs. The hold time value only impacts the alarm reporting time. To mark an interface up or down, you must also configure the physical interface hold time at the [edit interfaces *interface-name* hierarchy level].

Required Privilege Level

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

tti

Syntax tti *tti-identifier*;

Hierarchy Level [edit interfaces *interface-name* otn-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.
Statement introduced in Junos OS Release 13.3 for MX Series routers.

Description Specify trace identifier options.

Options *tti-identifier*—Trace identifier. It can be one of the following:

- **odu-dapi**—Optical Channel Data Unit (ODU) Destination Access Point Identifier.
- **odu-expected-receive-dapi**—ODU Expected Receive Destination Access Point Identifier.
- **odu-expected-receive-sapi**—ODU Expected Receive Source Access Point Identifier.
- **odu-sapi**—ODU Source Access Point Identifier.
- **otu-dapi**—Optical Channel Transport Unit (OTU) Destination Access Point Identifier.
- **otu-expected-receive-dapi**—OTU Expected Receive Destination Access Point Identifier.
- **otu-expected-receive-sapi**—OTU Expected Receive Source Access Point Identifier.
- **otu-sapi**—OTU Source Access Point Identifier.

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

tx-power

Syntax tx-power *dbm*;

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Transmit laser output power (dBm).

Default If you don’t specify a value, the default transmit laser output power is –2 dBm.

Options ***dbm***—Transmit power value.

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

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 465
- **optics-options on page 1288**
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475

warning

Syntax

```
warning low-light-warning {
   (link-down | syslog);
}
```

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information

Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold.

Options

link-down—Drop the 10-Gigabit Ethernet link and marks link as down.

syslog—Write the optics information to the system log.

Required Privilege Level

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

Related Documentation

- Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning on page 432
- **optics-options on page 1288**
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475

wavelength

Syntax wavelength *nm*;

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description For 10-Gigabit or 100-Gigabit Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.

Options *nm*—Wavelength value. It can be one of the following:

NOTE: **All values are displayed. However, if you configure a value that is not supported by the device, an error message is displayed and the device is not tuned to the specified wavelength.**

- **1528.38**—1528.38 nanometers (nm), corresponds to a 50-GHz grid
- **1528.77**—1528.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.16**—1529.16 nm, corresponds to a 50-GHz grid
- **1529.55**—1529.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.94**—1529.94 nm, corresponds to a 50-GHz grid
- **1530.33**—1530.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1530.72**—1530.72 nm, corresponds to a 50-GHz grid
- **1531.12**—1531.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1531.51**—1531.51 nm, corresponds to a 50-GHz grid
- **1531.90**—1531.90 nm, corresponds to 50-GHz and 100-GHz grids
- **1532.29**—1532.29 nm, corresponds to a 50-GHz grid
- **1532.68**—1532.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.07**—1533.07 nm, corresponds to a 50-GHz grid
- **1533.47**—1533.47 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.86**—1533.86 nm, corresponds to a 50-GHz grid
- **1534.25**—1534.25 nm, corresponds to 50-GHz and 100-GHz grids
- **1534.64**—1534.64 nm, corresponds to a 50-GHz grid

- **1535.04**—1535.04 nm, corresponds to 50-GHz and 100-GHz grids
- **1535.43**—1535.43 nm, corresponds to a 50-GHz grid
- **1535.82**—1535.82 nm, corresponds to 50-GHz and 100-GHz grids
- **1536.22**—1536.22 nm, corresponds to a 50-GHz grid
- **1536.61**—1536.61 nm, corresponds to 50-GHz and 100-GHz grids
- **1537.00**—1537.00 nm, corresponds to a 50-GHz grid
- **1537.40**—1537.40 nm, corresponds to 50-GHz and 100-GHz grids
- **1537.79**—1537.79 nm, corresponds to a 50-GHz grid
- **1538.19**—1538.19 nm, corresponds to 50-GHz and 100-GHz grids
- **1538.58**—1538.58 nm, corresponds to a 50-GHz grid
- **1538.98**—1538.98 nm, corresponds to 50-GHz and 100-GHz grids
- **1539.37**—1539.37 nm, corresponds to a 50-GHz grid
- **1539.77**—1539.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1540.16**—1540.16 nm, corresponds to a 50-GHz grid
- **1540.56**—1540.56 nm, corresponds to 50-GHz and 100-GHz grids
- **1540.95**—1540.95 nm, corresponds to a 50-GHz grid
- **1541.35**—1541.35 nm, corresponds to 50-GHz and 100-GHz grids
- **1541.75**—1541.75 nm, corresponds to a 50-GHz grid
- **1542.14**—1542.14 nm, corresponds to 50-GHz and 100-GHz grids
- **1542.54**—1542.54 nm, corresponds to a 50-GHz grid
- **1542.94**—1542.94 nm, corresponds to 50-GHz and 100-GHz grids
- **1543.33**—1543.33 nm, corresponds to a 50-GHz grid
- **1543.73**—1543.73 nm, corresponds to 50-GHz and 100-GHz grids
- **1544.13**—1544.13 nm, corresponds to a 50-GHz grid
- **1544.53**—1544.53 nm, corresponds to 50-GHz and 100-GHz grids
- **1544.92**—1544.92 nm, corresponds to a 50-GHz grid
- **1545.32**—1545.32 nm, corresponds to 50-GHz and 100-GHz grids
- **1545.72**—1545.72 nm, corresponds to a 50-GHz grid
- **1546.12**—1546.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1546.52**—1546.52 nm, corresponds to a 50-GHz grid
- **1546.92**—1546.92 nm, corresponds to 50-GHz and 100-GHz grids
- **1547.32**—1547.32 nm, corresponds to a 50-GHz grid
- **1547.72**—1547.72 nm, corresponds to 50-GHz and 100-GHz grids

- **1548.11**—1548.11 nm, corresponds to a 50-GHz grid
- **1548.51**—1548.51 nm, corresponds to 50-GHz and 100-GHz grids
- **1548.91**—1548.91 nm, corresponds to a 50-GHz grid
- **1549.32**—1549.32 nm, corresponds to 50-GHz and 100-GHz grids
- **1549.72**—1549.72 nm, corresponds to a 50-GHz grid
- **1550.12**—1550.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1550.52**—1550.52 nm, corresponds to a 50-GHz grid
- **1550.92**—1550.92 nm, corresponds to 50-GHz and 100-GHz grids
- **1551.32**—1551.32 nm, corresponds to a 50-GHz grid
- **1551.72**—1551.72 nm, corresponds to 50-GHz and 100-GHz grids
- **1552.12**—1552.12 nm, corresponds to a 50-GHz grid
- **1552.52**—1552.52 nm, corresponds to 50-GHz and 100-GHz grids
- **1552.93**—1552.93 nm, corresponds to a 50-GHz grid
- **1553.33**—1554.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1553.73**—1554.73 nm, corresponds to a 50-GHz grid
- **1554.13**—1554.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1554.54**—1554.54 nm, corresponds to a 50-GHz grid
- **1554.94**—1554.94 nm, corresponds to 50-GHz and 100-GHz grids
- **1555.34**—1555.34 nm, corresponds to a 50-GHz grid
- **1555.75**—1555.75 nm, corresponds to 50-GHz and 100-GHz grids
- **1556.15**—1556.15 nm, corresponds to a 50-GHz grid
- **1556.55**—1556.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1556.96**—1556.96 nm, corresponds to a 50-GHz grid
- **1557.36**—1557.36 nm, corresponds to 50-GHz and 100-GHz grids
- **1557.77**—1557.77 nm, corresponds to a 50-GHz grid
- **1558.17**—1558.17 nm, corresponds to 50-GHz and 100-GHz grids
- **1558.58**—1558.58 nm, corresponds to a 50-GHz grid
- **1558.98**—1558.98 nm, corresponds to 50-GHz and 100-GHz grids
- **1559.39**—1559.39 nm, corresponds to a 50-GHz grid
- **1559.79**—1559.79 nm, corresponds to 50-GHz and 100-GHz grids
- **1560.20**—1560.20 nm, corresponds to a 50-GHz grid
- **1560.61**—1560.61 nm, corresponds to 50-GHz and 100-GHz grids
- **1561.01**—1561.01 nm, corresponds to a 50-GHz grid

- **1561.42**—1561.42 nm, corresponds to 50-GHz and 100-GHz grids
- **1561.83**—1561.83 nm, corresponds to a 50-GHz grid
- **1562.23**—1562.23 nm, corresponds to 50-GHz and 100-GHz grids
- **1562.64**—1562.64 nm, corresponds to a 50-GHz grid
- **1563.05**—1563.05 nm, corresponds to 50-GHz and 100-GHz grids
- **1563.45**—1563.45 nm, corresponds to a 50-GHz grid
- **1563.86**—1563.86 nm, corresponds to 50-GHz and 100-GHz grids
- **1564.27**—1564.27 nm, corresponds to a 50-GHz grid
- **1564.68**—1564.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.09**—1565.09 nm, corresponds to a 50-GHz grid
- **1565.50**—1565.50 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.90**—1565.90 nm, corresponds to a 50-GHz grid
- **1566.31**—1566.31 nm, corresponds to 50-GHz and 100-GHz grids
- **1566.72**—1566.72 nm, corresponds to a 50-GHz grid
- **1567.13**—1567.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1567.54**—1567.54 nm, corresponds to a 50-GHz grid
- **1567.95**—1567.95 nm, corresponds to 50-GHz and 100-GHz grids
- **1568.36**—1568.36 nm, corresponds to a 50-GHz grid
- **1568.77**—1568.77 nm, corresponds to 50-GHz and 100-GHz grids

Default: 1550.12—1550.12 nm, corresponds to 50-GHz and 100-GHz grids

Required Privilege Level

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

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 465
- Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength on page 466
- **show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) on page 1803**

CHAPTER 38

Configuration Statements (OAM-CFM)

- short-name-format on page 1050
- sla-iterator-profile on page 1051
- sla-iterator-profiles on page 1052

action-profile (Applying to CFM)

Syntax action-profile *profile-name*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Identify the action profile to use.

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

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

action-profile (Defining for CFM)

Syntax

```
action-profile profile-name {
  event {
    ais-trigger-condition {
      adjacency-loss;
      all-defects;
      cross-connect-ccm;
      erroneous-ccm;
      receive-ais;
    }
    interface-status-tlv (down | lower-layer-down);
    port-status-tlv blocked;
    rdi;
  }
  action {
    interface-down;
    log-and-generate-ais {
      interval(1m | 1s);
      levelvalue;
      priority value;
    }
  }
  default-actions {
    interface-down;
  }
}
```

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

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure a name and default action for an action profile.

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

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- default-actions on page 1005
- event (CFM) on page 1012
- interface-down on page 1017

action-profile (MEP)

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

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Attach the configured action profile to the MEP depending on the hierarchy level.

Options ***action-profile-name***—Name of the action profile that is configured for the CFM MEP and the server MEP.

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

Related Documentation
- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

ais-trigger-condition

Syntax

```
ais-trigger-condition {
  adjacency-loss;
  all-defects;
  cross-connect-ccm;
  erroneous-ccm;
  receive-ais;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the defect conditions that generate an alarm indication signal (AIS).

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

all-defects

Syntax all-defects;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the defect condition that raises an alarm indication signal when any or all possible defects occur in the maintenance domain level.

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

auto-discovery

Syntax auto-discovery;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable the MEP to accept continuity check messages from all remote MEPs.

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

avg-fd-twoway-threshold

Syntax
avg-fd-twoway-threshold *avg-fd-twoway-threshold-value* {
flap-trap-monitor *seconds*;}

Hierarchy Level
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information
Statement introduced in Junos OS Release 15.1 for MX Series routers.

Description
Configure the threshold value for average frame delay, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay is exceeded, an SNMP trap is generated for ETH-DM. SNMP traps are triggered if you configure either the average frame-delay threshold or the average inter-frame delay variance threshold. If you do not configure either the frame-delay threshold or the frame delay variation threshold, no SNMP traps are generated. You can configure these threshold values only with a two-way ETH-DM SLA iterator.

Frame delay refers to the difference, in microseconds, between the time a frame is sent and when it is received. Frame delay variation refers to the difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called "frame jitter." For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

Options
avg-fd-twoway-threshold-value—Threshold value for average frame delay, in microseconds, for two-way ETH-DM.

Range: 1 through 4294967295 microseconds

flap-trap-monitor *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

Range: 1 through 360 seconds

Required Privilege Level
Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- **flap-trap-monitor on page 1013**
- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

avg-ifdv-twoway-threshold

Syntax

avg-ifdv-twoway-threshold *avg-ifdv-twoway-threshold-value*; {
flap-trap-monitor *seconds*; }

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information

Statement introduced in Junos OS Release 15.1 for MX Series routers.

Description

Configure the threshold value for average frame delay variation, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay variation is exceeded, an SNMP trap is generated for ETH-DM. SNMP traps are triggered if you configure either the average frame-delay threshold or the average inter-frame delay variance threshold. If you do not configure either the frame-delay threshold or the frame delay variation threshold, no SNMP traps are generated. You can configure these threshold values only with a two-way ETH-DM SLA iterator.

Frame delay refers to the difference, in microseconds, between the time a frame is sent and when it is received. Frame delay variation refers to the difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called "frame jitter." For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics. In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

Options

avg-ifdv-twoway-threshold-value—Threshold value for average frame delay variation, in microseconds, for two-way ETH-DM.

Range: 1 through 4294967295 microseconds

flap-trap-monitor *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

Range: 1 through 360 seconds

Required Privilege Level

configure—To enter configuration mode.
control—To modify any configuration.

Related Documentation

- flap-trap-monitor on page 1013
- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

avg-flr-forward-threshold

Syntax avg-flr-forward-threshold *avg-flr-forward-threshold-value*;{
flap-trap-monitor *seconds*; }

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX Series routers.

Description Configure the threshold value for average frame loss ratio, in milli-percent, in the upstream or forward direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average forward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. SNMP traps are triggered if you configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold. If you do not configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold, no SNMP traps are generated. You can configure these threshold values with an SLA iterator for ETH-SLM and ETH-LM.

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic.

Options ***avg-flr-forward-threshold-value***—Threshold value for average frame loss ratio in the forward or upstream direction, in milli-percent, for ETH-SLM and ETH-LM.

Range: 1 through 100000 milli-percent

flap-trap-monitor *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

Range: 1 through 360 seconds

Required Privilege Level configure—To enter configuration mode.
control—To modify any configuration.

Related Documentation

- **flap-trap-monitor on page 1013**
- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

avg-flr-backward-threshold

Syntax
avg-flr-backward-threshold *avg-flr-backward-threshold-value*;{
flap-trap-monitor *seconds*; }

Hierarchy Level
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information
Statement introduced in Junos OS Release 15.1 for MX Series routers.

Description
Configure the threshold value for average frame loss ratio, in milli-percent, in the backward or downstream direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average backward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. SNMP traps are triggered if you configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold. If you do not configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold, no SNMP traps are generated. You can configure these threshold values with an SLA iterator for ETH-SLM and ETH-LM.

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic.

Options
avg-flr-backward-threshold-value—Threshold value for average frame loss ratio in the backward or downstream direction, in milli-percent, for ETH-SLM and ETH-LM.

Range: 1 through 100000 milli-percent

flap-trap-monitor *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

Range: 1 through 360 seconds

Required Privilege Level
configure—To enter configuration mode.
control—To modify any configuration.

Related Documentation
- **flap-trap-monitor on page 1013**
- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

calculation-weight

Syntax

```
calculation-weight {
  delay delay-value;
  delay-variation delay-variation-value;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information

Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description

Configure the calculation weight for delay and delay variation.

NOTE: This option is applicable only for two-way delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- delay on page 1006
- delay-variation on page 1008

clear-action (CFM)

Syntax

```
clear-action {
  interface-down peer-interface;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name*]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Clear the action or actions to be taken when the connectivity fault management event occurs. You cannot configure multiple actions at this time. Only one action can be configured. This limitation affects both the **action** and **clear-action** statements.

Options

peer-interface—Name of the peer interface.

Required Privilege Level

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

Related Documentation

- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625

continuity-check

Syntax

```
continuity-check {
  convey-loss-threshold;
  hold-interval minutes;
  interface-status-tlv;
  interval (10m | 10s | 1m | 1s| 100ms | 10ms);
  loss-threshold number;
  port-status-tlv;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Specify continuity check protocol options.

Options

convey-loss-threshold—Enable loss-threshold-tlv transmission.

hold-interval *minutes*—Specify the continuity check hold-interval, in minutes.

interface-status-tlv—Enable interface-status-tlv transmission.

interval *(10m | 10s | 1m | 1s| 100ms | 10ms)*—Specify the continuity check interval.

loss-threshold *minutes*—Specify the loss-threshold, in minutes.

port-status-tlv—Enable port-status-tlv transmission.

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

Related Documentation

- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612

convey-loss-threshold

Syntax convey-loss-threshold;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Description Enable loss-threshold-tlv transmission.

During a unified ISSU, the control plane may go down for several seconds and cause CFM continuity check packets to get dropped. This may cause the remote maintenance endpoint (MEP) to detect a connectivity loss and mark the MEP as down. To keep the MEP active during a unified ISSU, the loss threshold TLV communicates the minimum threshold value the receiving MEP requires to keep the MEP active. The receiving MEP parses the TLV and updates the loss threshold value, but only if the new threshold value is greater than the locally configured threshold value. You can control the transmission of the loss threshold TLV in continuity check messages PDUs. The **convey-loss-threshold** statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress.

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

cross-connect-ccm

Syntax cross-connect-ccm;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the defect condition that raises an alarm indication signal when any cross-connect continuity check messages (CCMs) are received by the MEP.

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

cycle-time

Syntax cycle-time *cycle-time-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the time (in milliseconds) taken between back-to-back transmissions of SLA frames for a single connection.

Options ***cycle-time-value***—Cycle time value in milliseconds.

Range: 10 through 3,600,000

Default: 1000

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

data-tlv-size

Syntax data-tlv-size *size*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure the size of the data TLV portion of the Y.1731 data frame.

Options *size*—Size of the data TLV portion of the Y.1731 data frame.

NOTE: This option is applicable only for two-way delay measurement.

Range: 1 through 1400 bytes

Default: 1

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- sla-iterator-profile on page 1051
- Configuring a Remote MEP with an Iterator Profile on page 765

default-actions

Syntax

```
default-actions {
  interface-down;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Define the action to be taken when connectivity to the remote MEP is lost.

Default If no action is configured, no action is taken.

Options **interface-down**—When a remote MEP connectivity failure is detected, bring the interface down.

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

Related Documentation

- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625

delay

Syntax delay *delay-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name* calculation-weight]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the calculation weight for delay.

Options ***delay-value***—Calculation weight for delay.

NOTE: This option is applicable only for two-way delay measurement.

Range: 1 through 65,535

Default: 1

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- calculation-weight on page 999

delegate-server-processing

Syntax delegate-server-processing;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring]

Release Information Statement introduced in Junos OS Release 11.1.

Description For Ethernet interfaces on MX Series routers , enable server-side processing for two-way delay measurement and loss measurement.

By default, the processing is done by the Routing Engine.

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

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723

delay-variation

Syntax delay-variation *delay-variation-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name* calculation-weight]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the calculation weight for delay variation.

Options ***delay-variation-value***—Calculation weight for delay variation.

NOTE: This option is applicable only for two-way delay measurement.

Range: 1 through 65,535

Default: 1

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- calculation-weight on page 999

detect-loc

Syntax detect-loc;

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

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

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify whether Ethernet OAM continuity checks are performed for an individual remote maintenance end point (MEP).

When you configure the **detect-loc** statement at [**edit protocols oam ethernet link-fault-management interface** ***interface-name***] hierarchy level, a loss-of-continuity (LOC) defect is raised when the peer is not found within a period that is equal to 3 times the current keepalive pdu interval. When an LOC defect is raised, a syslog error message is generated .

NOTE: When you configure the **detect-loc** statement at the [**edit protocols oam ethernet link-fault-management interface** ***interface-name***] hierarchy level, any **action-profile** configured to bring down the interface is executed when an LOC defect is detected. However, the **action-profile** is not executed if you have not configured **detect-loc** statement the **detect-loc** statement at the [**edit protocols oam ethernet link-fault-management interface** ***interface-name***] hierarchy level.

To view the current LOC status of an interface, execute the **show oam ethernet link-fault-management** command.

Default The MEP does not generate LOC defect messages by default.

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- remote-mep on page 1048

direction

Syntax direction (up | down);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the direction of the MEP.

Options **up**—An UP MEP CCM is transmitted out of every logical interface which is part of the same bridging or vpls instance except for the interface configured on this MEP.

NOTE: The up direction for MEP is not supported on T Series routers.

down—Down MEP CCMs are transmitted only out the interface configured on this MEP.

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

enhanced-cfm-mode

Syntax enhanced-cfm-mode *enhanced-cfm-mode*;

Hierarchy Level [edit logical-systems *name* protocols oam ethernet connectivity-fault-management],
[edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 17.2R1 for MX Series routers.

Description Enables enhanced CFM mode. When you enable enhanced CFM mode, Junos OS supports 32, 000 maintenance association end points (MEPs) and maintenance intermediate points (MIPs) each per chassis for bridge, VPLS, L2VPN, and CCC domains. To support enhanced CFM mode, configure the network services mode on the router as **enhanced-ip**.

NOTE: After enabling CFM mode, restart CFM for the changes to take effect. If you do not restart CFM, CFM automatically restarts after 1 minute.

Required Privilege Level Routing

Related Documentation

- Enabling Enhanced Connectivity Fault Management Mode on page 666

erroneous-ccm

Syntax erroneous-ccm;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the defect condition that raises an alarm indication signal when any cross-connect continuity check messages (CCMs) with an unexpected MEP ID or an erroneous maintenance domain level are received by the MEP.

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

event (CFM)

Syntax

```
 event {
   adjacency-loss;
   interface-status-tlv [lower-layer-down down];
   port-status-tlv blocked;
   rdi;
 }
}
```

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

Release Information Statement introduced in Junos OS Release 10.1

Description Configure threshold values for connectivity fault management events in an action profile.

Options

adjacency-loss—Connectivity is lost.

interface-status-tlv [lower-layer-down down]—Values that need to be monitored in interface status TLV.

port-status-tlv—Values that need to be monitored in port status TLV.

rdi—RDI received from some MEP.

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

Related Documentation

- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- interface-status-tlv on page 1018
- port-status-tlv on page 1041

flap-trap-monitor

Syntax flap-trap-monitor *seconds*;

Hierarchy Level [edit logical-systems *name* protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
[edit logical-systems *name* protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name* avg-fd-twoway-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles*profile-name* avg-fd-twoway-threshold]
[edit logical-systems *name* protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name* avg-ifdv-twoway-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles*profile-name* avg-ifdv-twoway-threshold]
[edit logical-systems *name* protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name* avg-flr-forward-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles*profile-name*avg-flr-forward-threshold]
[edit logical-systems *name* protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name* avg-flr-backward-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles*profile-name*avg-flr-backward-threshold]

Release Information Statement introduced in Junos OS Release 17.2R1 for MX Series routers.

Description Enables damping of jnxSoamPmThresholdCrossingAlarm traps sent to the network management system (NMS) by summarizing the flap occurrences over a period of time and sends a single jnxSoamPmThresholdFlapAlarm notification to the NMS. You can enable damping at the global level for the iterator or you can enable damping at the individual threshold type of an iterator. You can specify the duration of time for summarizing flap occurrences.

Options *seconds*—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the NMS.

Range: 1 through 360 seconds

Required Privilege Level routing

Related Documentation

- avg-fd-twoway-threshold on page 995
- avg-ifdv-twoway-threshold on page 996
- avg-flr-backward-threshold on page 998
- avg-flr-forward-threshold on page 997
- sla-iterator-profiles on page 1052

- Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS on page 766

hardware-assisted-timestamping

Syntax hardware-assisted-timestamping;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring]

Release Information Statement introduced in Junos OS Release 9.5.

Description For Ethernet interfaces on Enhanced and Enhanced Queuing Dense Port Concentrators (DPCs) in MX Series routers only, enable hardware-assisted timestamping support for Ethernet frame delay measurement.

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

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

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768
- Enabling the Hardware-Assisted Timestamping Option on page 779

hardware-assisted-keepalives

Syntax hardware-assisted-keepalives [enable | disable];

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring]

Release Information Statement introduced in Junos OS Release 14.2R1.

Description For Ethernet interfaces on Modular Port Concentrators (MPCs) in MX Series routers only, delegate the transmission of the continuity check messages (CCMs) to the forwarding ASIC (that is, to the hardware) by enabling inline transmission of CCMs. Inline transmission of CCMs is also known as Inline-KA.

By default, CCMs are transmitted by the CPU of the MPC and not by the hardware. If the duration between transmissions of CCMs is low or if the CCMs for a specific line card scale, we recommend that you enable delegation of the transmission of CCMs to the hardware. By enabling inline transmission of CCMs, you can achieve maximum scaling of CCMs.

NOTE: The supported CCM interval values for MPC7E and MPC2E NG Q (MPC2E-3D-NG-Q) are 100ms, 10ms, and 1s.

Default Inline transmission is disabled by default.

Options **enable**—Enable inline transmission of CCMs.

NOTE: Inline transmission of CCMs is not enabled when there is a CFM session already established. To enable inline transmission, you must first deactivate the CFM session using the **deactivate** command and then reactivate the CFM session using the **activate** command.

disable—Disable inline transmission of CCMs.

NOTE: After disabling inline transmission of CCMs, you must reboot the router for the changes to take effect.

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

Related Documentation

- Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling on page 795
- Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades on page 800

hold-interval (OAM)

Syntax hold-interval *minutes*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description The time to wait in minutes before flushing the maintenance association end point (MEP) database, if no updates occur. The configurable range is 1 minute through 30240 minutes.The default value is 10 minutes.

NOTE: Hold timer based flushing is applicable only for auto discovered remote MEPs and not for statically configured remote MEPs.

Options *minutes*—Time to wait, in minutes.

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

Related Documentation

- Continuity Check Protocol Parameters Overview on page 611
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612

instance

Syntax instance *vpls-instance-name*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the VPLS instance of the default maintenance domain.

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

Related Documentation

- Configuring Maintenance Intermediate Points (MIPs) on page 604
- maintenance-domain on page 1030

interface-down

Syntax interface-down;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* default-actions]

Release Information Statement introduced in Junos OS Release 8.5.

Description Bring the interface down when a remote MEP connectivity failure is detected.

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

Related Documentation

- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625

interface-status-tlv

Syntax interface-status-tlv [down lower-layer-down];

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 9.6.

Description Defines an **action-profile** consisting of various events and the action. Based on values of **interface-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using **action-profile** options.

Options **down**—When the incoming CCM packet contains interface status TLV with value down, the action will be triggered for this action-profile.

lower-layer-down—When the incoming CCM packet contains interface status TLV with value lower-layer-down, the action will be triggered for this action-profile.

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

Related Documentation

- Configuring Remote MEP Action Profile Support on page 645

interface-status-send-rdi

Syntax interface-status-send-rdi *interface-status-send-rdi*;

Hierarchy Level [edit logical-systems *name* protocols oam ethernet connectivity-fault-management maintenance-domain *name* maintenance-association *name* continuity-check],
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *name* maintenance-association *name* continuity-check]

Release Information Statement introduced in Junos OS Release 17.3R1 for MX Series Routers.

Description Configure CFM to propagate the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages when the interface is down. When the status of the EVPN provider edge device is standby, the EVPN VPWS service is notified and it sets the interface status to CCC-down. When the interface status is CCC-down, it indicates that the provider edge service is down. When you enable CFM monitoring, CFM propagates the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages. Thus, the customer edge device is aware that the provider edge device is down.

Usually, when the interface goes down, CFM propagates the status of the provider edge device via the interface status TLV. If the customer edge device does not support the interface status TLV, you can use the RDI bit to propagate the status of the provider edge device.

Required Privilege Level routing—To view this statement in the configuration.

Related Documentation

- interface-status-tlv on page 1018
- Configuring Port Status TLV and Interface Status TLV on page 635
- Understanding CFM Monitoring between CE and PE Devices on page 667

interval

Syntax interval (100ms | 10m | 10ms | 10s | 1m | 1s);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 8.4.
Option **10ms** introduced in Junos OS Release 9.1.
Third-party interoperability during a unified in-service software upgrade (ISSU) introduced in Junos OS Release 17.1.

Description Configure the interval between successive transmissions of continuity check messages (CCMs) as part of the connectivity fault detection strategy. When the receiving maintenance association end point (MEP) does not receive a CCM at the configured interval, the **loss-threshold** statement determines how many CCMs can be lost before the sending MEP is marked as down. The **hold-interval** statement then determines the frequency at which the database of MEPs in the maintenance association (MA) is flushed in the absence of updates.

During a unified in-service software upgrade (ISSU), Junos OS connectivity fault management (CFM) works when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit CCM (continuity check message) PDUs during the unified ISSU upgrade. For this feature to work, you must enable Packet Forwarding Engine keepalives with the **hardware-assisted-keepalives** statement, and configure the interval between CCMs to be 1 second with **interval** statement.

NOTE: For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can disable PPM only on the Packet Forwarding Engine. To disable PPM on the Packet Forwarding Engine, use the no-delegate-processing statement at the [edit routing-options ppm] hierarchy level.

NOTE: A continuity check interval of 10 milliseconds is not supported for CFM sessions over a label-switched interface (LSI).

Options **100ms**—100 milliseconds.

10m—10 minutes.

10ms—10 milliseconds.

10s—10 seconds.

1m—1 minute.

1s—1 second.

Default: 1m

Required Privilege Level

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

Related Documentation

- Continuity Check Protocol Parameters Overview on page 611
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612
- Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades on page 800

interval (CFM MEP)

Syntax

interval (1m | 1s);

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action log-and-generate-ais]

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Configure the interval between AIS messages that are to be received by the MEP as either 1 minute or 1 second.

Required Privilege Level

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

iteration-count

Syntax iteration-count *count-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure the number of iterations for which the connection partakes in the iterator for acquiring SLA measurements.

Options ***count-value***—Number of iterations for which the connection should partake in the iterator for acquiring SLA measurements.

Range: 1 through 65,535

Default: 0 (or infinite iterations)

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- sla-iterator-profile on page 1051
- Configuring a Remote MEP with an Iterator Profile on page 765

iteration-period

Syntax iteration-period *iteration-period-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure the iteration period, which is the maximum number of cycles per iteration (that is, the number of connections registered to an iterator cannot exceed this value).

Options ***iteration-period-value***—Maximum number of cycles per iteration.

Range: 1 through 2000

Default: 2000

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

level

Syntax

level *number*;

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in junos os release 12.1X48 for PTX Series Packet Transport Routers.

Description

A number used in connectivity fault management (CFM) messages to identify the maintenance association. The number is embedded in each of the CFM frames. CFM messages within a given level are processed by maintenance end points (MEPs) at the same level. For example, the operator domain can be level 0, the provider domain can be level 3, and the customer domain can be level 7.

Options

number—A number used to identify the maintenance domain to which the CFM message belongs.

Range: 0 through 7

Required Privilege Level

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

Related Documentation

- Creating a Maintenance Domain on page 603

level (CFM MEP)

Syntax level *value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action log-and-generate-ais]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the server maintenance domain level for the MEP.

Options **value**—Maintenance domain level.

Range: 1 through 7

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

Related Documentation
- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

linktrace

Syntax

```
linktrace {
  age (30m | 10m | 1m | 30s | 10s);
  path-database-size path-database-size;
}
```

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

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure connectivity fault management linktrace parameters.

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

Related Documentation
- Configuring Linktrace Protocol in CFM on page 626

log-and-generate-ais

Syntax

```
log-and-generate-ais {
  interval (1m | 1s);
  level level;
  priority level;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action]

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Configure the action be taken when an AIS alarm is detected. The action includes generating and logging the AIS statistics along with the interval between AIS messages, the server maintenance domain level, and the priority of the AIS message.

Options

The other 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

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

loss-threshold

Syntax loss-threshold *number*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Specify the number of continuity check messages lost before marking the remote MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

Options *number*—The number of continuity check messages that can be lost before the remote MEP is considered down.

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

Related Documentation

- Continuity Check Protocol Parameters Overview on page 611
- Configuring Continuity Check Protocol Parameters for Fault Detection on page 612

lowest-priority-defect

Syntax lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-err-xcon | xcon)

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify the lowest priority defect that is allowed to generate a Fault Alarm whenever CFM detects a defect. This configuration is done at the MEP level.

Options Specify one of the following lowest priority defect options:

all-defects—Allows all defects.

err-xcon—Allows only erroneous CCM and cross-connect CCM defects.

mac-rem-err-xcon—Allows only MAC, not receiving CCM, erroneous CCM, and cross-connect defects.

no-defect—Allows no defects.

rem-err-xcon—Allows only not receiving CCM, erroneous CCM, and cross-connect CCM defects.

xcon—Allows only cross-connect CCM defects.

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

Related Documentation

- *Configuring the Maintenance End Point Lowest Priority Defect*

maintenance-association

Syntax

```
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 {
    hold-interval minutes;
    interval (10m | 10s | 1m | 1s| 100ms);
    loss-threshold number;
  }
  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;
      }
    }
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description

Configure the name of the maintenance association in IEEE-compliant format.

Options

ma-name—The name of the maintenance association within the maintenance domain.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Creating a Maintenance Association on page 610
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

maintenance-domain

Syntax

```
maintenance-domain domain-name {
  bridge-domain name <vlan-id [ vlan-ids ]>;
  instance vpls-instance-name;
  level number;
  maintenance-association ma-name {
    protect-maintenance-association  protect-ma-name;
    remote-maintenance-association  remote-ma-name;
    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    continuity-check {
      hold-interval minutes;
      interval (10m | 10s | 1m | 1s| 100ms);
      loss-threshold number
    }
    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;
        }
      }
    }
    mip-half-function(none | default | explicit);
    name-format (character-string | none | dns | mac+2oct);
  }
  virtual-switch name {
    bridge-domain name <vlan-id [ vlan-ids ]>;
  }
}
```

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

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Support for multiple down MEP introduced in Junos OS Release 15.1R1 for MX Series Routers.

Description Configure the name of the maintenance domain in IEEE-compliant format.

NOTE: For MX Series Routers, you can configure multiple down MEPs for a single instance of maintenance domain identifier and maintenance association name to monitor services provided on Virtual Private LAN Service (VPLS), bridge, circuit cross-connect (CCC), and IPv4 domains.

Options *domain-name*—Name of the maintenance domain.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Creating a Maintenance Domain on page 603
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

measurement-interval

Syntax measurement-interval (5|15|30|60)

Hierarchy Level [edit protocols oam ethernet cfm performance-monitoring]

[edit protocols oam ethernet cfm performance-monitoring sla-iterator-profiles *profile-name*]

Release Information

Description Configure measurement interval to be used for a performance monitoring session. You must configure the measurement-interval at the **[edit protocols oam ethernet cfm performance-monitoring]** hierarchy level, which is a global level parameter. You can override the configured value by specifying a measurement-interval for the iterator profile at the **[edit protocols oam ethernet cfm performance-monitoring sla-taerator-profiles *profile-name*]** hierarchy level.

NOTE: When you configure when MEF-36-compliant performance monitoring, you must also configure an **enhanced-sla-iterator** at the **[edit protocols oam ethernet cfm performance-monitoring]** hierarchy level.

Default 15 minutes

Required Privilege Level configure—To enter configuration mode.
control—To modify any configuration.

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596
- Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36 on page 601

measurement-type

Syntax measurement-type (loss | statistical-loss-measurement | two-way-delay);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 11.1.
The **statistical-loss-measurement** option introduced in Junos OS Release 11.2.

Description Configure the measurement type for the service level agreement (SLA) frames. An SLA frame is a type of packet used to measure frame loss in Ethernet connections.

Options **loss**—Use Y.1731-compliant line module (LM) frames to measure frame loss.

statistical-loss-measurement— Use Y.1731-compliant two-way data module (DM) frames to statistically measure frame loss.

two-way-delay—Use Y.1731-compliant two-way DM frames to measure frame loss.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754

mep

Syntax

```
mep mep-id {
  action-profile action-profile-name
  auto-discovery;
  direction (up | down);
  interface interface-name (protect | working);
  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;
    }
  }
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

The numeric identifier of the maintenance association end point (MEP) within the maintenance association.

Options

mep *mep-id*—Specify the numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811
- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

mip-half-function

Syntax

mip-half-function (none | default | explicit);

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-managementmaintenance-domain *md-name*],
[edit protocols oam ethernet connectivity-fault-managementmaintenance-association *ma-name*]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Specify the OAM Ethernet CFM maintenance domain MIP half functions.

NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations are the same.

Options

none—Specify to not use the mip-half-function.

default—Specify to use the default mip-half-function.

explicit—Specify an explicit mip-half-function.

Required Privilege Level

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

Related Documentation

- Creating a Maintenance Domain on page 603
- maintenance-domain on page 1030

name-format

Syntax name-format (character-string | none | dns | mac+2oct);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Specify the format of the maintenance domain name.

Options **character-string**—The name is an ASCII character string.

none—The maintenance domain name is not used.

dns—The name is in domain name service (DNS) format. For example: www.juniper.net.

mac+2oct—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.

Default: character-string

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

Related Documentation

- Creating a Maintenance Association on page 610
- Creating a Maintenance Domain on page 603

path-database-size

Syntax path-database-size *path-database-size*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management linktrace]

Release Information Statement introduced in Junos OS Release 8.5.

Description Number of linktrace reply entries to be stored per linktrace request.

Options **path-database-size**—Database size.

Range: 1 through 255

Default: 64

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

Related Documentation

- Configuring Linktrace Protocol in CFM on page 626

performance-monitoring

Syntax

```
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);
    }
  }
}
```

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

Release Information Statement introduced in Junos OS Release 9.5.

Description Specify performance monitoring support for Ethernet frame delay measurement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 723
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 768
- Enabling the Hardware-Assisted Timestamping Option on page 779

policer (CFM Global)

Syntax

```
policer {
  all cfm-policer-name;
  continuity-check cfm-policer-name;
  other cfm-policer-name;
  }
}
```

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

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a policer at the global level to police the CFM traffic belonging to all sessions.

Options

continuity-check *cfm-policer-name*—Police all continuity check packets with the policer specified.

other *cfm-policer-name*—Police all non-continuity check packets with the policer specified.

all *cfm-policer-name*—Police all CFM packets with policer specified. If the **all** option is used, then you cannot specify above two options.

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

Related Documentation

- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- **policer (CFM Session) on page 1040**

policer (CFM Session)

Syntax

```
policer {
  all cfm-policer-name;
  continuity-check cfm-policer-name;
  other cfm-policer-name;
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *name* level *number* maintenance-association *name*]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a separate policer to rate-limit packets specific to that session.

Options

- **continuity-check** ***cfm-policer-name***—Police continuity check packets belonging to this session.

  other ***cfm-policer-name***—Police all non-continuity check packets belonging to this session.

  all ***cfm-policer-name***—Police all CFM packets belonging to this session. If the **all** option is used, then you cannot specify the above two options.

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

Related Documentation

- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- policer (CFM Global) on page 1039

port-status-tlv

Syntax port-status-tlv blocked;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *tlv-action* event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]

Release Information Statement introduced in Junos OS Release 9.6.

Description Define an **action-profile** consisting of various events and the action. Based on values of **port-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using action-profile options.

Options **blocked**—When the incoming CCM packet contains port status TLV with value blocked, the action will be triggered for this action-profile.

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

Related Documentation

- Configuring a CFM Action Profile to Specify CFM Actions for CFM Events on page 625
- Configuring Remote MEP Action Profile Support on page 645

priority (Protocols OAM)

Syntax priority *priority-value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-namesla-iterator-profile*]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure the priority of the iterator profile, which is the **vlan-pcp** value that is sent in the Y.1731 data frames.

Options *priority-value*—Priority value, which is the **vlan-pcp** value that is sent in the Y.1731 data frames.

Range: 0 through 7

Default: 0

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- sla-iterator-profile on page 1051
- Configuring a Remote MEP with an Iterator Profile on page 765

priority (CFM MEP)

Syntax priority *value*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action log-generate-ais]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the 802.1p priority of the AIS packet.

Options ***value***—Priority level.

Range: 0 through 7

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

priority (OAM Connectivity-Fault Management)

Syntax

priority *number*;

Hierarchy Level

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

For EX Series Switches:

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

Release Information

Statement introduced in Junos OS Release 8.4.

Description

IEEE 802.1p priority bits used by the continuity check messages.

Options

number—Configure the IEEE 802.1p priority bits to be used in the VLAN header of the CFM packets.

Range: 0 through 7

Required Privilege Level

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

protocol (Server MEP)

Syntax

```
protocol (l2circuit | l2vpn | ethernet) {
  interface interface-name;
}
```

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

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the protocol as Layer 2 circuit, Layer 2 VPN, or Ethernet and associate the interface to the protocol that needs to be monitored for ETH-AIS.

Options

l2circuit—Configure the protocol for the server MEP as Layer 2 circuit.

l2vpn—Configure the protocol for the server MEP as Layer 2 VPN.

ethernet—Configure the protocol for the server MEP as Ethernet.

interface *interface-name*—The interface that is to be associated with the protocol that needs monitoring for ETH-AIS.

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

protect-maintenance-association (OAM)

Syntax protect-maintenance-association *protect-ma-name*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 11.4

Description Configure the name of the protect transport path for the maintenance-association.

Options ***protect-ma-name***—The name of the protect transport path.

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

receive-ais

Syntax receive-ais;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]

Release Information Statement introduced in Junos OS Release 14.2.

Description Send a message to the peer MEPs when an AIS message is received by a peer MEP at its own maintenance level.

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

remote-maintenance-association (OAM)

Syntax remote-maintenance-association *remote-ma-name*;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the name of the remote maintenance association.

Options ***remote-ma-name***—Name of the remote maintenance association.

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

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613

remote-mep

Syntax

```
remote-mep mep-id {
  action-profile profile-name;
  sla-iterator-profile profile-name {
    data-tlv-size size;
    iteration-count count-value;
    priority priority-value;
  }
  detect-loc;
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

Options *mep-id*—Numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring a MEP to Generate and Respond to CFM Protocol Messages on page 613
- detect-loc on page 1009

sendid-tlv

Syntax

```
sendid-tlv {
  send-chassis-tlv;
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name* maintenance-association *maintenance-association-name* continuity-check]

Release Information Statement introduced in Junos OS Release 16.1R2.

Description Configures Junos OS to send the sender ID TLV along with the packets. The sender ID TLV is an optional TLV that is sent in continuity check messages (CCMs), loopback messages, and Link Trace Messages (LTMs), as specified in the IEEE 802.1ag standard. TLVs (type, length, and value) are described in the IEEE 802.1ag standard for Connectivity Fault Management (CFM) as a method of encoding variable-length and optional information in a protocol data unit (PDU).

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

Related Documentation

- Junos OS Support for Chassis ID TLV on page 602

short-name-format

Syntax short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Specify the name format of the maintenance association name.

Options **character-string**—The name is an ASCII character string.

vlan—The primary VLAN identifier.

2octet—A number in the range 0 through 65,535.

rfc-2685-vpn-id—A VPN identifier that complies with RFC 2685.

Default: character-string

NOTE: The PTX Series Packet Transport Routers support the **vlan** and **2octet** options only.

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

Related Documentation

- Creating a Maintenance Association on page 610

sla-iterator-profile

Syntax

```
sla-iterator-profile profile-name {
  data-tlv-size size;
  iteration-count count-value;
  priority priority-value;
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 11.1.

Description

Configure a remote MEP with an iterator profile and specify the options.

Options

profile-name—Name of the iterator profile configured for a remote MEP. For more information about configuring a remote MEP with an iterator profile, see "Configuring a Remote MEP with an Iterator Profile" on page 765.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Verifying the Configuration of an Iterator Profile on page 756
- Managing Iterator Statistics on page 759
- **sla-iterator-profiles on page 1052**

sla-iterator-profiles

Syntax

```
sla-iterator-profiles {
  profile-name {
    avg-fd-twoway-threshold;
      avg-ifdv-twoway-threshold;
      avg-flr-forward-threshold;
      avg-flr-backward-threshold;
      calculation-weight {
      delay delay-weight;
      delay-variation delay-variation-weight;
    }
    cycle-time milliseconds;
    flap-trap-monitor seconds
    iteration-period iteration-period-value;
    measurement-type (loss | statistical-frame-loss | two-way-delay);
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management performance-monitoring]

Release Information Statement introduced in Junos OS Release 11.1.

Description Configure an iterator application and specify the iterator profile options.

Options *profile-name*—Name of the iterator profile. For more information about configuring the iterator profile, see "Configuring an Iterator Profile" on page 754.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 754
- Configuring a Remote MEP with an Iterator Profile on page 765
- Verifying the Configuration of an Iterator Profile on page 756
- Managing Iterator Statistics on page 759

CHAPTER 39

Configuration Statements

802.3ad

Syntax

```
802.3ad {
  primary | backup;
  ae interface-number ;
  lacp {
    port-priority priority-number;
  }
  link-index index-number
  distribution-list distribution-list-number
}
```

Hierarchy Level

[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
primary and **backup** options added in Junos OS Release 8.3.

Description

Specify aggregated Ethernet logical interface number.

Options

bundle—Join an aggregated Ethernet interface.

ae *interface-number*—Aggregated Ethernet logical interface number. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. On MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.

primary | backup—For link protection configurations, specify the link as primary link or backup link for egress traffic.

lacp—Configure Link Aggregation Control Protocol. Specify the port priority in the range 0 through 65535. Default port-priority is 127.

link-index—Specify the desired child link index within the aggregated Ethernet Interface. Index number of the logical interface reflects its initialization sequence.

distribution-list—For targeted distribution, specify the distribution list to which the interface belongs.

Required Privilege Level

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

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 108
- Configuring Aggregated Ethernet Link Protection on page 148

accept-source-mac

Syntax

```
accept-source-mac {
  mac-address mac-address {
    policer {
      input cos-policer-name;
      output cos-policer-name;
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet intelligent queuing (IQ) interfaces only, accept traffic from and to the specified remote media access control (MAC) address.

The **accept-source-mac** statement is equivalent to the **source-address-filter** statement, which is valid for aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only. To allow the interface to receive packets from specific MAC addresses, include the **accept-source-mac** statement.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

The remaining statements are explained separately. See CLI Explorer.

NOTE: The policer statement is not supported on PTX Series Packet Transport Routers.

NOTE: On QFX platforms, if you configure source MAC addresses for an interface using the *static-mac* or *persistent-learning* statements and later configure a different MAC address for the same interface using the **accept-source-mac** statement, the MAC addresses that you previously configured for the interface remain in the ethernet-switching table and can still be used to send packets to the interface.

Required Privilege Level

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

Related Documentation

- Configuring MAC Address Filtering on page 544
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 16
- **source-filtering on page 1366**

access-concentrator

Syntax access-concentrator *name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information Statement introduced before Junos OS Release 7.4.
Support at the **[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** and **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]** hierarchy levels introduced in Junos OS Release 10.1.
Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description Configure an alternative access concentrator name in the AC-NAME tag in a PPPoE control packet for use with a dynamic PPPoE subscriber interface. If you do not configure the access concentrator name, the AC-NAME tag contains the system name.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options *name*—Name of the access concentrator.

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

Related Documentation

- Identifying the Access Concentrator on page 356
- *Configuring the PPPoE Family for an Underlying Interface*
- *Configuring Dynamic PPPoE Subscriber Interfaces*
- PPPoE Overview on page 350

account-layer2-overhead (PIC Level)

Syntax account-layer2-overhead;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Enable the automatic adjustment of Layer 2 overhead in bytes, which is the octet adjustment per packet, based on the encapsulation on the logical interface for the total octet count for ingress and egress traffic on all the interfaces in the PIC.

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

Related Documentation

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics on page 533
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 536
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 537

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 700

action-profile

List of Syntax

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 1068
Syntax: EX Series Switches and NFX Series Devices on page 1068

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches

```
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;
  }
}
```

Syntax: EX Series Switches and NFX Series Devices

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

Hierarchy Level

[edit protocols oam ethernet link-fault-management]

Release Information

Statement introduced in Junos OS Release 8.5 for T, M, MX and ACX Series Routers, SRX Series Firewalls, and EX Series Switches, .
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description

Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.

The remaining statements are explained separately. See CLI Explorer.

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

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring an OAM Action Profile on page 698
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

adaptive

Syntax

```
adaptive {
  pps;
  scan-interval multiple;
  tolerance tolerance-percentage;
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options load-balance]

Release Information Statement introduced in Junos OS Release 13.2R3 for MX Series Routers.
Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 15.1X53-D10 for the QFX Series.

Description Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.

Options **pps**—(PTX Series only) The type of traffic rate among the members of the AE bundle is measured packets per second. The default rate type is bytes per second.

scan-interval *multiple*—(PTX Series only) Scan interval, as a multiple of a 30-second interval.

Range: 1 through 5

Default: 1

tolerance *tolerance-percentage*—(MX Series and PTX Series) Limit to the variance in the packet traffic flow to the aggregated Ethernet links in a percentage.

Range: 1 through 100 percent

Default: 20 percent

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

Related Documentation

- Understanding Aggregated Ethernet Load Balancing on page 158
- Example: Configuring Aggregated Ethernet Load Balancing on page 163

address

Syntax

```
address address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  master-only;
  multipoint-destination address dlci dlci-identifier;
  multipoint-destination address {
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
        length);
      queue-length number;
    }
    vci vpi-identifier.vci-identifier;
  }
  primary;
  preferred;
  virtual-gateway-address
  (vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise–interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
    priority-number number;
    track {
      priority-cost seconds;
      priority-hold-time interface-name {
        interface priority;
        bandwidth-threshold bits-per-second {
          priority;
        }
      }
      route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*],

[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

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

Description Configure the interface address.

Options *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

NOTE: If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, and the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/8;
     }
  }
}
```

For more information on configuring the same address on multiple interfaces, see *Configuring the Interface Address*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

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

Related Documentation

- *Configuring the Protocol Family*
- *Junos OS Administration Library*
- *family*
- negotiate-address on page 1270
- *unnumbered-address (Ethernet)*

adjacency-loss

Syntax adjacency-loss;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* event ais-trigger-condition]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure the defect condition that raises an alarm indication signal when physical connectivity is lost between peer MEPs.

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

Related Documentation

- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805
- Configuring ETH-AIS on a CFM MEP on page 811

advertisement-interval

Syntax advertisement-interval *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6 for MX Series and T Series routers.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.

Description Configure an interval for LLDP advertisement.

For switches configured for Link Layer Discovery Protocol, configure the frequency at which LLDP advertisements are sent.

The **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value, or an error will be returned when you attempt to commit the configuration.

NOTE: The default value of **transmit-delay** is 2 seconds. If you configure the **advertisement-interval** as less than 8 seconds and you do not configure a value for **transmit-delay**, the default value of **transmit-delay** is automatically changed to 1 second in order to satisfy the requirement that the **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value.

Default Disabled.

Options *seconds*—Interval between LLDP advertisement.

Default: 30

Range: 5 through 32768

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

Related Documentation

- Configuring LLDP on page 338
- *show lldp*
- *Configuring LLDP (CLI Procedure)*
- *Understanding LLDP and LLDP-MED on EX Series Switches*

- *transmit-delay*
- *Understanding LLDP*

age

Syntax age (30m | 10m | 1m | 30s | 10s);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management linktrace]

Release Information Statement introduced in Junos OS Release 8.5.

Description Time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.

Default 10 minutes

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

Related Documentation

- Configuring Linktrace Protocol in CFM on page 626

agent-specifier

Syntax

```
agent-specifier {
  aci circuit-id-string ari remote-id-string {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    static-interface interface-name;
  }
}
```

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*]

Release Information Statement introduced in Junos OS Release 10.0.
drop, **delay**, **terminate**, **dynamic-profile**, **routing-instance**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Specify the action taken by the interface for the specified agent circuit identifier/agent remote identifier (ACI/ARI) pair when the interface receives a PPPoE Active Discovery Initiation (PADI) control packet that includes the vendor-specific tag with ACI/ARI pair information. You can configure an ACI/ARI pair for a named service, **empty** service, or **any** service in a PPPoE service name table. A maximum of 8000 ACI/ARI pairs are supported per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the named, **empty**, and **any** service entries in the service name table.

You can use an asterisk (*) as a wildcard character to match ACI/ARI pairs, the ACI alone, or the ARI alone. The asterisk can be placed only at the beginning, the end, or both the beginning and end of the identifier string. You can also specify an asterisk alone for either the ACI or the ARI. You cannot specify only an asterisk for both the ACI and the ARI. When you specify a single asterisk as the identifier, that identifier is ignored in the PADI packet.

For example, suppose you care about matching only the ACI and do not care what value the ARI has in the PADI packet, or even whether the packet contains an ARI value. In this case you can set the ***remote-id-string*** to a single asterisk. Then the interface ignores the ARI received in the packet and the interface takes action based only on matching the specified ACI.

Default The default action is terminate.

Options **aci *circuit-id-string***—Identifier for the agent circuit ID that corresponds to the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

ari *remote-id-string*—Identifier for the subscriber associated with the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information*

aggregate (Gigabit Ethernet CoS Policer)

Syntax

```
aggregate {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- *premium (Hierarchical Policer)*
- ieee802.1p on page 1188

aggregated-devices

Syntax

```
aggregated-devices {
  ethernet {
    device-count number;
    lacp {
      link-protection {
        non-revertive;
      }
      system-priority;
    }
  }
  sonet {
    device-count number;
    }
  maximum-links maximum-links-limit;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 7.4.
Support for LACP link protection and system priority introduced in Junos OS Release 9.3.

Description Configure properties for aggregated devices on the router.

Options 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 Junos OS for Supporting Aggregated Devices on page 129

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  ethernet-switch-profile {
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
    (mac-learn-enable | no-mac-learn-enable);
  }
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    link-protection {
      disable;
    (revertive | non-revertive);
    periodic interval;
    system-priority priority;
    system-id system-id;
  }
  link-protection;
  load-balance;
  link-speed speed;
  logical-interface-chassis-redundancy;
  logical-interface-fpc-redundancy;
  (loopback | no-loopback);
  minimum-links number;
  rebalance-periodic time hour:minute <interval hours>;
  source-address-filter {
    mac-address;
    (source-filtering | no-source-filtering);
  }
}
```


Hierarchy Level [edit interfaces ae*x*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure aggregated Ethernet-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Ethernet Interfaces Overview on page 3

alarms

Syntax alarms;

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced in JUNOS Release 10.1.

Description For 10-Gigabit Ethernet DPCs, configure the DPC to drop the interface link when the receive power falls below the alarm threshold.

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

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 465

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 706

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 704

arp (Interfaces)

Syntax

arp *ip-address* (mac | multicast-mac) *mac-address* publish;

```
arp {
  aging-timer minutes;
  gratuitous-arp-delayseconds;
  gratuitous-arp-on-ifup;
  interfaces {
    interface-name {
      aging-timer minutes;
    }
  }
  passive-learning;
  purging;
}
```

Syntax (EX Series)

```
arp {
  aging-timer minutes;
}
```

Hierarchy Level

[edit system]

[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*],

[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

Release Information

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

Description

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

NOTE: For EX-Series switches, set only the time interval between ARP updates.

Options

ip-address—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing **address** statement.

mac *mac-address*—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn*:*nn*:*nn*:*nn*:*nn*:*nn***. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

multicast-mac *mac-address*—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn*:*nn*:*nn*:*nn*:*nn*:*nn***. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.

NOTE: For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

aging-timer—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

passive-learning (QFX-Series only)—Configure backup VRRP routers or switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the master router fails, the backup router must learn all entries present in the ARP cache of the master router. Configuring passive learning reduces transition delay when the backup router is activated.

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

Related Documentation
- Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses on page 372
- *Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Junos OS System Basics Configuration Guide* .

asynchronous-notification

Syntax
(asynchronous-notification | no-asynchronous-notification);

Hierarchy Level
[edit interfaces ge-*fpc/pic/port* gigether-options]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description
(MX Series routers, T Series routers) For all Gigabit Ethernet interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit), configure support for notification of link down alarm generation and transfer.

(M120 and M320 routers) For all 10-Gigabit Ethernet PIC interfaces, configure support for notification of link down alarm generation and transfer.

- **asynchronous-notification**—Support notification of link down alarm generation and transfer.
- **no-asynchronous-notification**—Prohibit notification of link down alarm generation and transfer.

Default
Support for notification of link down alarm generation and transfer is not enabled.

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

Related Documentation
- Gigabit Ethernet Notification of Link Down Alarm Overview on page 431
- Configuring Gigabit Ethernet Notification of Link Down Alarm on page 432

authentication-access-control (MX Series in Enhanced LAN Mode)

Syntax

```
authentication-access-control {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag ;
    }
  uac-policy;
    authentication-profile-name access-profile-name;
    no-mac-table-binding {
        interface  interface-names
        static  mac-address
    }
    static  mac-address {
        interface interface-names;
        vlan-assignment (vlan-id |vlan-name);
    }
    interface (all | [ interface-names ]) {
        session-expiry seconds;
        quiet-period seconds;
        reauthentication {
            interval seconds;
        }
        retries number;
        server-timeout seconds;
        supplicant (single | single-secure | multiple);
        dot1x {
            disable;
            guest-vlan (vlan-id | vlan-name);
            mac-radius {
                flap-on-disconnect;
                restrict;
            }
            maximum-requests number;
            no-reauthentication;
            server-fail (deny | permit | use-cache | vlan-id | vlan-name);
            server-reject-vlan (vlan-id | vlan-name) {
                eapol-block;
                block-interval block-interval;
            }
            supplicant-timeout seconds;
            transmit-period seconds;
        }
    }
  }
  (captive-portal | no-captive-portal);
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure an authenticator for 802.1X and captive-portal authentication.

The remaining statements are explained separately. See CLI Explorer.

NOTE: You cannot configure 802.1X user authentication on interfaces that have been enabled for Q-in-Q tunneling.

Default No static MAC address or VLAN is configured.

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

authentication-profile-name

Syntax authentication-profile-name *access-profile-name*;

Hierarchy Level [edit protocols dot1x authenticator]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the RADIUS authentication profile to use for user authentication when establishing an IEEE 802.1x Port-Based Network Access Control (**dot1x**) connection.

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authenticator on page 1087**
- **dot1x on page 1124**

authenticator

Syntax

```
authenticator {
  authentication-profile-name access-profile-name;
  interface interface-id {
    maximum-requests integer;
    quiet-period seconds;
    reauthentication (disable | interval seconds);
    retries integer;
    server-timeout seconds;
    supplicant (single);
    supplicant-timeout seconds;
    transmit-period seconds;
  }
}
```

Hierarchy Level [edit protocols dot1x]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify an authentication profile for user or client authentication and configure the Ethernet interface for 802.1x protocol operation.

Options **authentication-profile-name** ***access-profile-name***—Specifies the RADIUS authentication profile for user or client authentication.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authentication-profile-name on page 1086**
- **dot1x on page 1124**

auto-negotiation

Syntax (auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-offline)>;

Hierarchy Level [edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. For EX Series switches, explicitly enable autonegotiation only.

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

NOTE: On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

NOTE: On ACX Series routers, when you configure fiber interfaces (fiber media mode) to operate at 1 Gbps, you need to always enable autonegotiation (**auto-negotiation**) to negotiate the speed and duplex settings. You cannot disable autonegotiation (**no-auto-negotiation**) in the fiber media mode. In copper interfaces (copper media mode), autonegotiation is enabled by default. To disable autonegotiation, you need to explicitly configure the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options **remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

Default: local-interface-online

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

Related Documentation

- Gigabit Ethernet Autonegotiation Overview on page 551
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

auto-reconnect

Syntax auto-reconnect *seconds*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces *interface-name*
unit *logical-unit-number* pppoe-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description PPP over Ethernet interfaces, configure the amount of time to wait before reconnecting after a session has terminated.

Options ***seconds***—Time to wait before reconnecting after a session has terminated.

Range: 0 through 4,294,967,295 seconds

Default: 0 (immediately)

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

Related Documentation

- Configuring the PPPoE Automatic Reconnect Wait Timer on page 356
- *Junos OS Interfaces and Routing Configuration Guide*

bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

Syntax

bandwidth-limit *bps*;

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* aggregate],
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* premium]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define a policer to apply to nonpremium traffic.

Options

bps—Bandwidth limit, in bits per second. Specify either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 32 Kbps through 32 gigabits per second (Gbps). For IQ2 and IQ2-E interfaces 65,536 bps through 1 Gbps. For 10-Gigabit IQ2 and IQ2-E interfaces 65,536 bps through 10 Gbps.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- **burst-size-limit (Policer for Gigabit Ethernet Interfaces) on page 1097**

bridge-domain

Syntax

bridge-domain *name*;
  vlan-id [*vlan-identifiers*];
}

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name*],
[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name* virtual-switch *virtual-switch-name*]

Release Information

Statement introduced in Junos OS Release 9.4.

Description

(MX Series routers only) Specify the OAM Ethernet CFM maintenance domain bridge domain.

Options

name—Specify the name of the bridge domain.

vlan-identifiers—Specify one or more VLAN identifiers.

Required Privilege Level

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

Related Documentation

- Configuring Maintenance Intermediate Points (MIPs) on page 604
- **maintenance-domain on page 1030**

bridge-domains

Syntax

```
bridge-domains {
  bridge-domain-name {
    bridge-options {
      ...bridge-options-configuration...
    }
    domain-type bridge;
    interface interface-name;
    no-irb-layer-2-copy;
    no-local-switching;
    routing-interface routing-interface-name;
    vlan-id (all | none | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer number inner number;
    bridge-options {
      interface interface-name {
        mac-pinning
        static-mac mac-address;
      }
      interface-mac-limit limit;
      mac-statistics;
      mac-table-size limit;
      no-mac-learning;
    }
  }
}
```

Hierarchy Level

[edit],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.
Support for the **no-irb-layer-2-copy** statement added in Junos OS Release 10.2.

Description

(MX Series routers only) Configure a domain that includes a set of logical ports that share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Options

bridge-domain-name—Name of the bridge domain.

NOTE: You cannot use the slash (/) character as part of the bridge domain name. If you do, the configuration will not commit.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring a Bridge Domain*
- *Configuring a Layer 2 Virtual Switch*

bfd-liveness-detection (LAG)

Syntax

bfd-liveness-detection {
  authentication {
    algorithm *algorithm-name*;
    key-chain *key-chain-name*;
    loose-check;
  }
  detection-time {
    threshold *milliseconds*;
  }
  holddown-interval *milliseconds*;
  local-address *bfd-local-address*;
  minimum-interval *milliseconds*;
  minimum-receive-interval *milliseconds*;
  multiplier *number*;
  neighbor *bfd-neighbor-address*;
  no-adaptation;
  transmit-interval {
    minimum-interval *milliseconds*;
    threshold *milliseconds*;
  }
  version (1 | automatic);
}

Hierarchy Level **[edit interfaces aex aggregated-ether-options]**

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure Bidirectional Forwarding Detection (BFD) timers and authentication for aggregated Ethernet interfaces.

Options

holddown-interval *milliseconds*— Specify a time limit, in milliseconds, indicating the time that a BFD session remains up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address *bfd-local-address*— Specify the loopback address or the AE interface address of the source of the BFD session.

NOTE: Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD **local-address** against the interface or loopback IP address before the configuration commit. Junos OS performs this check on both IPv4 and IPv6 micro BFD address configurations, and if they do not match, the commit fails.

minimum-interval *milliseconds*— Specify a minimum time interval after which the local routing device transmits a BFD packet and then expects to receive a reply from the BFD neighbor. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** statement.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*— Specify the minimum time interval after which the routing device expects to receive a reply from the BFD neighbor.

Range: 1 through 255,000

multiplier *number*— Specify the number of BFD packets that were not received by the BFD neighbor before the originating interface is declared down.

Range: 1 through 255

neighbor *bfd-neighbor-address*— Specify the loopback address or the AE interface address of a remote destination to send BFD packets.

no-adaptation— Disable the BFD adaptation. Include this statement if you do not want the BFD sessions to adapt to changing network conditions. We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

version— Configure the BFD version to detect (BFD version 1) or autodetect (the BFD version).

NOTE: **The version option is not supported on the QFX Series.**

Default: automatic

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *authentication*
- *detection-time*
- *transmit-interval*
- *Configuring Independent Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG on page 186
- Understanding Independent Micro BFD Sessions for LAG on page 183

burst-size-limit (Policer for Gigabit Ethernet Interfaces)

Syntax burst-size-limit *bytes*;

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* aggregate],
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* premium]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a policer to apply to nonpremium traffic.

Options *bytes*—Burst length.

Range: 1500 through 100,000,000 bytes

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- bandwidth-limit (Policer for Gigabit Ethernet Interfaces) on page 1091

cak (MX Series)

Syntax cak *hexadecimal-number*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* pre-shared-key]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers. Statement introduced in Junos OS Release 17.3R1 for MX10003 3D Universal Edge Routers.

Description Specifies the connectivity association key (CAK) for a pre-shared key.

A pre-shared key includes a connectivity association key name (CKN) and a CAK. A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link

Default No CAK exists, by default.

Options ***hexadecimal-number***—The key name, in hexadecimal format.

The key name is 32 hexadecimal characters in length. If you enter a key name that is less then 32 characters long, the remaining characters are set to 0.

On MX10003 router, to maximize the security, it is recommended to configure CAK of even length.

- If you configure CAK of length that is less than 32 hexadecimal digits and if cipher-suite is gcm-aes-128/gcm-aes-256 and less than 64 hexadecimal digits, then the following warning message is displayed:**warning: To maximize security, recommend configuring all 32 digits of pre-shared-key cak** or **warning: To maximize security, recommend configuring all 64 digits of pre-shared-key cak**
- On MX10003 router, if you configure the length of CAK to an odd value, then the following warning message is displayed: **To maximize security, it is recommended to configure pre-shared-key cak of even length**

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

captive-portal (MX Series in Enhanced LAN Mode)

Syntax (captive-portal | no-captive-portal);

Hierarchy Level [edit protocols authentication-access-control]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Specify whether captive portal authentication needs to be enabled or disabled. You can set up captive portal authentication (hereafter referred to as captive portal) on a switch to redirect Web browser requests to a login page that requires the user to input a username and password. Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

Default Not enabled

Options **captive-portal**—Enable captive portal authentication.

no-captive-portal—Disable captive portal authentication.

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

captive-portal-custom-options (MX Series in Enhanced LAN Mode)

Syntax

captive-portal-custom-options {
  banner-message *string*;
  footer-bgcolor *color*;
  footer-message *string*;
  footer-text-color *color*;
  form-header-bgcolor *color*;
  form-header-message *string*;
  form-header-text-color *color*;
  form-reset-label *label name*;
  form-submit-label *label name*;
  header-bgcolor *color*;
  header-logo *filename*;
  header-message *string*;
  header-text-color *color*;
  post-authentication-url *url-string*;
}

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Specify the design elements of a captive portal login page.

Options

banner-message—The first screen displayed before the captive portal login page is displayed—for example, a disclaimer message.

Range: 1–2047 characters

footer-bgcolor —The hexadecimal color code for the color of the footer bar across the bottom of the captive portal login page—for example, #2E8B57 (sea green).

Values: # symbol followed by six characters.

footer-message—Text message displayed in the footer bar across the bottom of the captive portal login page.

Range: 1–2047 characters

Default: **Copyright @2010, Juniper Networks Inc.**

footer-text-color — Color of the text in the footer.

Default: The default color is white.

form-header-bgcolor —The hexadecimal color code for the background color of the header bar across the top of the form area of the captive portal login page.

Values: # symbol followed by six characters.

form-header-message—Text message displayed in the header bar across the top of the form area of the captive portal login page.

Range: 1–255 characters

Default: **Captive Portal User Authentication**

form-header-text-color—Color of the text in the form header.

Default: The default color is black.

form-reset-label—Label displayed in the button that the user can select to clear the username and password fields on the form.

Range: 1–255 characters

Default: **Reset**

form-submit-label —Label displayed in the button that the user selects to submit their login information—for example, **Log In** .

Range: 1–255 characters

Default: **Log In**

header-bgcolor—The hexadecimal color code for the color of the header bar across the top of the captive portal login page.

Values: # symbol followed by six characters.

header-logo—Filename of the file containing the image of the logo displayed at the top of the captive portal login page. The image file can be in GIF, JPEG, or PNG format.

Default: The Juniper Networks logo

header-message—Text displayed in the header bar across the bottom of the captive portal login page.

Range: 1–2047 characters

Default: **User Authentication**

header-text-color—Color of the text in the header.

Default: The default color is white.

post-authentication-url—URL to which the users are directed upon successful authentication—for example **www.mycafe.com**.

Range: 1–255 characters

Default: The page originally requested by the user.

Required Privilege Level

routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

centralized

Syntax centralized;

Hierarchy Level [edit protocols lacp ppm]

Release Information Statement introduced in Junos OS Release 9.4 for MX Series routers.
Statement introduced in Junos OS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.

This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** statement in the **[edit routing-options ppm]** hierarchy.

BEST PRACTICE: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

Default Distributed PPM processing is enabled for all packets that use PPM.

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

Related Documentation

- *Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- Configuring Distributed Periodic Packet Management on page 197
- *Configuring Link Aggregation*

cipher-suite (MACsec)

Syntax

cipher-suite *encryption-algorithm-name*;

Hierarchy Level

[edit security macsec connectivity-association *connectivity-association-name*]

Release Information

Statement introduced in Junos OS Release 16.2R1 for MX240, MX480, MX960, MX2020, and MX2010 routers.
Statement introduced in Junos OS Release 17.2R1 for QFX Series switches.
Statement introduced in Junos OS Release 17.3R1 for JNP-MIC1-MACSEC MIC on MX10003 routers.
Statement introduced in Junos OS Release 18.2R1 for EX Series switches.

Description

Specify the set of ciphers used to encrypt traffic on an Ethernet link that is secured with Media Access Control Security (MACsec). The encryption used by MACsec ensures that the data in the Ethernet frame cannot be viewed by anybody monitoring traffic on the link. MACsec encryption is optional and user-configurable. The configured cipher suites should be the same between MACsec peers.

MACsec utilizes the Galois/Counter Mode Advanced Encryption Standard (GCM-AES). The default cipher suite used for MACsec is GCM-AES-128, with a maximum key length of 128 bits. MACsec also supports GCM-AES-256, with a maximum key length of 256 bits.

GCM–AES–128 and GCM–AES–256 use a 32-bit packet number as part of the initial value that has to be unique for every packet sent with a given secure association key (SAK). When the permutations of the 32-bit packet number are exhausted, the SAK much be refreshed. The frequency of SAK refreshes can be reduced by using a cipher suite with Extended Packet Numbering (XPN), which increases the size of the packet number to 64-bits. Both GCM-AES-128 and GCM-AES-256 are available with XPN.

NOTE: When enabling MACsec on a QFX10016 or QFX10008 switch, we recommend using either the GCM-AES-XPN-128 or GCM-AES-XPN-256 cipher suite.

Default

If the **cipher-suite** statement is not configured, the default cipher suite used for encryption is GCM-AES-128.

Options

gcm-aes-128—GCM-AES-128 has a maximum key size of 128 bits.

gcm-aes-xpn-128—GCM-AES-XPN-128 has a maximum key size of 128 bits and extended packet number.

gcm-aes-256—GCM-AES-256 has a maximum key size of 256 bits.

gcm-aes-xpn-256—GCM-AES-XPN-256 has a maximum key size of 256 bits and extended packet number.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring Media Access Control Security (MACsec)*
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

ckn (MX Series)

Syntax ckn *hexadecimal-number*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* pre-shared-key]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.
Statement introduced in Junos OS Release 17.3R1 for MX10003 3D Universal Edge Routers.

Description Specifies the connectivity association key name (CKN) for a pre-shared key.

A pre-shared key includes a CKN and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link

Default No CKN exists, by default.

Options *hexadecimal-number*—The key name, in hexadecimal format.

The key name is 32 hexadecimal characters in length. If you enter a key name that is less then 32 characters long, the remaining characters are set to 0.

- On MX10003 router, if you configure the length of CKN to the value less than 64 hexadecimal digits, then the following warning message is displayed:

  warning: To maximize security, recommend configuring all 64 digits of pre-shared-key ckn

- On MX10003 router, if you configure the length of CKN to an odd value, then the commit will not be successful and the following error message is displayed:

  error: ckn: 'abcde': Must be an even-length string up to 64 hexadecimal digits (0-9, a-f, A-F)

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

classifier

Syntax

```
classifier {
  per-unit-scheduler {
    forwarding-class class-name {
      loss-priority (high | low);
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the classifier for the output priority map to be applied to outgoing frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Specifying an Output Priority Map on page 541
- input-priority-map on page 1200

clear

Syntax request protection-group ethernet-aps clear md <md> ma <ma>

Hierarchy Level [edit protocols protection-group ethernet-aps]

Description Clears the lockout, force switch, manual switch, exercise, and wait-to-restore (WTR) states.

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

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 213

client

Syntax client;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the router to operate in the PPPoE client mode.

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

Related Documentation

- Configuring the PPPoE Client Mode on page 357

community-vlans (MX Series)

Syntax community-vlans [*number number-number*];

Hierarchy Level [edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* ,
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers.

Description Configure the specified community VLAN to be a secondary VLAN of the specified primary VLAN. A *community* VLAN is used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN.

NOTE: When you specify this configuration statement, the VLAN ID of a logical interface that you associate with a bridge domain that matches with the VLAN ID or list of IDs that you specify using the **community-vlans** state is treated as a community port.

Options *number*—Individual VLAN IDs separated by a space.

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

Related Documentation

compatibility-version

Syntax compatibility-version;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify the compatible version mode to be used. When compatibility-version is set to value 1, the node operates in ITU-T Recommendation G.8032/Y.1344 version 1 compatible mode. In this mode all the supported external commands are blocked, ring-id is forced to be 1 and mode of operation is set to revertive mode.

Options

- **1**—Use ITU-T Recommendation G.8032/Y.1344 compatibile mode version 1.
- **2**—Use ITU-T Recommendation G.8032/Y.1344 compatibile mode version 2.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

connectivity-association (MACsec Interfaces for MX Series)

Syntax connectivity-association *connectivity-association-name*;

Hierarchy Level [edit security macsec interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Applies a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface.

Default No connectivity associations are associated with any interfaces.

Options ***connectivity-association-name***—Name of the MACsec connectivity association.

Range: 1 through 32 alphanumeric characters. Allowed characters are [a-z, A-Z, 0-9]

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

connectivity-association (MX Series)

Syntax

```
connectivity-association connectivity-association-name {
  exclude-protocol protocol-name;
  include-sci;
  mka {
    must-secure;
    key-server-priority priority-number;
    transmit-interval interval;
  }
  no-encryption;
  offset (0|30|50);
  pre-shared-key {
    cak hexadecimal-number;
    ckn hexadecimal-number;
  }
  replay-protect{
    replay-window-size number-of-packets;
  }
  secure-channel secure-channel-name {
    direction (inbound | outbound);
    encryption ;
    id {
      mac-address mac-address;
      port-id port-id-number;
    }
    offset (0|30|50);
    security-association security-association-number {
      key key-string;
    }
  }
  security-mode security-mode;
}
```

Hierarchy Level [edit security macsec]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Create or configure a MACsec connectivity association.

A connectivity association is not applying MACsec to traffic until it is associated with an interface. MACsec connectivity associations are associated with interfaces using the **interfaces** statement in the **[edit security macsec]** hierarchy.

Default No connectivity associations are present, by default.

Options *connectivity-association-name*—Name of the MACsec connectivity association.

Range: 1 through 32 alphanumeric characters. Allowed characters are [a-z, A-Z, 0-9]

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 Media Access Control Security (MACsec) on MX Series Routers*

connectivity-fault-management

Syntax

```
connectivity-fault-management {
  action-profile profile-name {
    action {
      interface-down;
      log-and-generate-ais {
        interval(1m | 1s);
        levelvalue;
        priority value;
      }
    }
    default-actions {
      interface-down;
    }
    event {
      ais-trigger-condition {
        adjacency-loss;
        all-defects;
        cross-connect-ccm;
        erroneous-ccm;
        receive-ais;
      }
      adjacency-loss;
      interface-status-tlv (down | lower-layer-down);
      port-status-tlv blocked;
      rdi;
    }
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    bridge-domain <vlan-id [ vlan-ids ]>;
    instance routing-instance-name;
    interface interface-name;
    level number;
    name-format (character-string | none | dns | mac+2oct);
    maintenance-association ma-name {
      protect-maintenance-association protect-ma-name;
      remote-maintenance-association remote-ma-name;
      short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
      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;
                    detect-loc;
                  }
                }
              }
            }
            virtual-switch routing-instance-name {
              bridge-domain name <vlan-ids [ vlan-ids ]>;
            }
          }
          no-aggregate-delegate-processing;
          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);
              }
            }
          }
        }
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 8.4.

Description For Ethernet interfaces on M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M120, M320, MX Series, and T Series routers, specify connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.
In Junos OS Release 9.3 and later, this statement is also supported on aggregated Ethernet interfaces.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 596

control-channel

Syntax
```
control-channel channel-name {
  vlan vlan-id;
  interface name interface-name
}
```

Hierarchy Level [edit protocols protection-group ethernet-ring *name* (east-interface | west-interface)]

Release Information
Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.

Options
vlan *vlan-id*—If the control channel logical interface is a trunk port, then a dedicated **vlan *vlan-id*** defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the ***vlan-id*** when the control channel logical interface is the trunk port.

interface name *interface-name*—Interface name of the control channel.

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

Related Documentation
- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

data-channel

Syntax

```
data-channel {
  vlan number;
}
```

Hierarchy Level

[edit protocols protection-group ethernet-ring *ring-name*]

Release Information

Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description

For Ethernet ring protection, configure a data channel to define a set of VLAN IDs that belong to a ring instance.

VLANs specified in the data channel use the same topology used by the ERPS PDU in the control channel. Therefore, if a ring interface is blocked in the control channel, all traffic in the data channel is also blocked on that interface.

Options

vlan ***number***—Specify (by VLAN ID) one or more VLANs that belong to a ring instance.

Required Privilege Level

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

Related Documentation

- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857
- Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers on page 864
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

delay (PPPoE Service Name Tables)

Syntax delay *seconds*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description Configure the PPPoE underlying interface on the router to wait a specified number of seconds after receiving a PPPoE Active Discovery Initiation (PADI) control packet from a PPPoE client before sending a PPPoE Active Discovery Offer (PADO) packet to indicate that it can service the client request

The router (PPPoE server) does not check whether another server has already sent a PADO packet during the delay period in response to the PPPoE client's PADI packet. It is up to the PPPoE client to determine whether another PPPoE server has responded to its PADI request, or if it must respond to the delayed PADO packet to establish a PPPoE session.

Options *seconds*—Number of seconds that the PPPoE underlying interface waits after receiving a PADI packet from a PPPoE client before sending a PADO packet in response.

Range: 1 through 120 seconds

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

Related Documentation
- *Configuring PPPoE Service Name Tables*

destination (IPCP)

Syntax destination *address* destination-profile *profile-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.

Options ***address***—IP address of the remote interface.

The remaining statement is explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring IPCP Options for Interfaces with PPP Encapsulation*
- address on page 1071
- negotiate-address on page 1270
- *Junos OS Administration Library*

device-count

Syntax device-count *number*;

Hierarchy Level [edit chassis aggregated-devices ethernet]
[edit chassis aggregated-devices sonet]

Release Information Statement introduced before Junos OS Release 7.4.
Statement functionality updated in Junos OS Release 14.2, as described below.

Description Configure the number of aggregated logical devices available to the router.

Starting in Junos release 14.2, for MX series routers, aggregate Ethernet interfaces created under a logical system can be individually named. Prior to 14.2, ae interfaces were named automatically (AE1, AE2) etc. upon setting the device count. This change allows administrators to use custom naming schemes. System resources are only allocated for named ae interfaces, regardless of how many were declared in the device count. (In Junos 14.2 and earlier, ae naming occurred automatically up to the number specified for device count, and system resources were allocated whether a given ae interface was used or not.)

Options *number*—Set the number of aggregated logical devices that will be available for configuration.

NOTE: **Starting with Junos OS Release 13.2, a maximum of 64 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces. In releases before Junos OS Release 13.2, a maximum of 16 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces.**

Range: 1 - 496. The upper limit for this value is system specific.

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

Related Documentation

- Configuring Junos OS for Supporting Aggregated Devices on page 129
- *Configuring Aggregated SONET/SDH Interfaces*

direction (MX Series)

Syntax direction (inbound | outbound);

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Configure whether the secure channel applies MACsec security to traffic entering or leaving an interface.

If you need to apply MACsec on traffic entering and leaving an interface, you need to create one secure channel to apply MACsec on incoming traffic and another secure channel to apply MACsec on outgoing traffic within the same connectivity association. When you associate the connectivity association with an interface, MACsec is applied on traffic entering and leaving that interface.

You only use this configuration option when you are configuring MACsec using static secure association keys (SAK) security mode. When you are configuring MACsec using static connectivity association keys (CAK) security mode, two secure channels that are not user-configurable—one inbound secure channel and one outbound secure channel—are automatically created within the connectivity association.

Default This statement does not have a default value.

If you have configured a secure channel to enable MACsec using static SAK security mode, you must specify whether the secure channel applies MACsec to traffic entering or leaving an interface. A candidate configuration that contains a secure channel that has not configured a direction cannot be committed.

Options **inbound**—Enable MACsec security on traffic entering the interface that has applied the secure channel.

outbound—Enable MACsec security on traffic leaving the interface that has applied the secure channel.

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 Media Access Control Security (MACsec) on MX Series Routers*

disable

Syntax disable;

Hierarchy Level [edit protocols lldp],
[edit protocols lldp interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp interface (all | *interface-name*)]

Release Information Statement introduced in Junos OS Release 9.6.

Description Disable LLDP globally or on an interface.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

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

Related Documentation

- Configuring LLDP on page 338

disable (Link Protection)

Syntax disable;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Disable LACP link protection on the interface.

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

Related Documentation

- *Configuring LACP for Aggregated Ethernet Interfaces*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

disable (802.1X for MX Series in Enhanced LAN Mode)

Syntax disable;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Disable 802.1X authentication on a specified interface or all interfaces.

Default 802.1X authentication is disabled on all interfaces.

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

distribution-list

Syntax distribution-list *distribution-list-number*;

Hierarchy Level [edit interfaces *interface name* gigether-options 802.3ad]
[edit dynamic-profiles *name* interfaces *name* gigether-options 802.3ad]
[edit dynamic-profiles *name* logical-systems *name* interfaces *name* gigether-options 802.3ad]

Release Information Statement introduced in Junos OS Release 16.1R1.

Description Specify a distribution list to a Gigabit Ethernet interface to carry traffic. You can then configure the distribution list as a primary list or a backup list for the members of an aggregated Ethernet bundle.

Example:

```
[edit]
user@router# set interfaces ge-0/0/3 gigether-options 802.3ad distribution-list dl1
```

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

Related Documentation

- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links on page 202
- targeted-options on page 1393
- targeted-distribution on page 1392

dot1p-priority

Syntax dot1p-priority *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify the IEEE 802.1p priority to be used in the transmitted RAPS protocol data units.

Options *number*—802.1p priority number.

Range: 0 through 7

Default: 0

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

dot1x

Syntax

```
dot1x {
  authenticator {
    authentication-profile-name access-profile-name;
    interface interface-id {
      maximum-requests integer;
      quiet-period seconds;
      reauthentication (disable | interval seconds);
      retries integer;
      server-timeout seconds;
      supplicant (single);
      supplicant-timeout seconds;
      transmit-period seconds;
    }
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.3.

Description For the MX Series only, specifies settings for using 802.1x Port-Based Network Access Control.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- authenticator on page 1087
- authentication-profile-name on page 1086
- interface (IEEE 802.1x) on page 1205

dot1x (MX Series in Enhanced LAN Mode)

Syntax

```
dot1x {
  disable;
  guest-vlan (vlan-id | vlan-name);
  mac-radius {
    flap-on-disconnect;
    restrict;
  }
  maximum-requests number;
  no-reauthentication;
  server-fail (deny | permit | use-cache | vlan-id | vlan-name);
  server-reject-vlan (vlan-id | vlan-name) {
    eapol-block;
    block-interval block-interval;
  }
  supplicant-timeout seconds;
  transmit-period seconds;
}
```

Hierarchy Level

[edit protocols authentication-access-control interface (all | [*interface-names*])]

Release Information

Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description

Configure 802.1X authentication for Port-Based Network Access Control. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANs).

The remaining statements are explained separately. See CLI Explorer.

Default

802.1X is disabled.

Required Privilege Level

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

domain-id

Syntax domain-id *domain-id*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3)],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3)]

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

Description Specify a domain ID for a route. The domain ID identifies the OSPF domain from which the route originated.

Options ***domain-id***—You can specify either an IP address or an IP address and a local identifier using the following format: ***ip-address*:*local-identifier***. If you do not specify a local identifier with the IP address, the identifier is assumed to have a value of 0.

Default: If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.

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

Related Documentation

- *Configuring an OSPF Domain ID*

drop (PPPoE Service Name Tables)

Syntax
drop;

Hierarchy Level
[edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information
Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description
Direct the router to drop (ignore) a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client that contains the specified service name tag or agent circuit identifier/agent remote identifier (ACI/ARI) information. This action effectively denies the client's request to provide the specified service, or to accept requests from the subscriber or subscribers represented by the ACI/ARI information.

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

Related Documentation
- *Configuring PPPoE Service Name Tables*

dynamic-profile (PPPoE Service Name Tables)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify a dynamic profile to instantiate a dynamic PPPoE interface. You can associate a dynamic profile with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The dynamic profile associated with a service entry in a PPPoE service name table overrides the dynamic profile associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **dynamic-profile** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **dynamic-profile** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options ***profile-name***—Name of the dynamic profile that the router uses to instantiate a dynamic PPPoE interface.

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

east-interface

Syntax

```
east-interface {
  node-id mac-address;
  control-channel channel-name {
    vlan number;
    interface name interface-name
  }
  interface-none
  ring-protection-link-end;
}
```

Hierarchy Level

[edit protocols protection-group ethernet-ring *ring-name*]

Release Information

Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description

Define one of the two interface ports for Ethernet ring protection, the other being defined by the **west-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

EX Series switches do not use the node-id statement--the node ID is automatically configured on the switches using the MAC address.

NOTE: Always configure this port first, before configuring the **west-interface** statement.

NOTE: The Node ID is not configurable on EX Series switches. The node ID is automatically configured using the MAC address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857
- west-interface on page 1452
- ethernet-ring on page 1153

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

egress-policer-overhead

Syntax egress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1.

Description Add the specified number of bytes to the actual length of an Ethernet frame when determining the actions of Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card. You can configure egress policer overhead to account for egress *shaping* overhead bytes added to output traffic on the line card.

On M Series and T Series routers, this statement is supported on Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs. On MX Series routers, this statement is supported for interfaces configured on Dense Port Concentrators (DPCs).

NOTE: **This statement is not supported on Modular Interface Cards (MICs) or Modular Port Concentrators (MPCs) in MX Series routers.**

Options *bytes*—Number of bytes added to a packet exiting an interface.

Range: 0–255 bytes

Default: 0

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

Related Documentation

- *egress-shaping-overhead*
- *Policer Overhead to Account for Rate Shaping Overview*
- *Example: Configuring Policer Overhead to Account for Rate Shaping*
- *Configuring a Policer Overhead*
- *CoS on Enhanced IQ2 PICs Overview*

encapsulation (Logical Interface)

Syntax encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-tcc-vc-mux | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | ethernet-ccc | ethernet-vpls | ethernet-vpls-fr | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-ppp | frame-relay-tcc | gre-fragmentation | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls | vxlan);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces rlsq *number* unit *logical-unit-number*]
[edit protocols evpn]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**ethernet**,**vlan-ccc**, and **vlan-tcc** options only).
Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers. Only the **atm-ccc-cell-relay** and **atm-ccc-vc-mux** options are supported on ACX Series routers.
Statement introduced in Junos OS Release 17.3R1 for QFX10000 Series switches (**ethernet-ccc** and **vlan-ccc** options only).

Description Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Options **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

NOTE: **The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.**

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

vxlan—Use VXLAN data plane encapsulation for EVPN.

Required Privilege Level

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

Related Documentation

- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring the Encapsulation for Layer 2 Switching TCCs*
- *Configuring Interface Encapsulation on Logical Interfaces*
- *Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects*
- *Circuit and Translational Cross-Connects Overview*
- Identifying the Access Concentrator on page 356
- *Configuring ATM Interface Encapsulation*
- Configuring VLAN and Extended VLAN Encapsulation on page 256
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM for Subscriber Access*
- *Understanding CoS on ATM IMA Pseudowire Interfaces Overview*
- *Configuring Policing on an ATM IMA Pseudowire*

encapsulation

List of Syntax

Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series on page 1135
Syntax for Logical Interfaces: SRX Series on page 1135

Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc |
  ethernet-bridge | ethernet-ccc | ethernet-over-atm | ethernet-tcc | ethernet-vpls |
  ethernet-vpls-fr | ether-vpls-over-atm-llc | ethernet-vpls-ppp | extended-frame-relay-ccc
  | extended-frame-relay-ether-type-tcc | extended-frame-relay-tcc | extended-vlan-bridge
  | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services
  | flexible-frame-relay | frame-relay | frame-relay-ccc | frame-relay-ether-type |
  frame-relay-ether-type-tcc | frame-relay-port-ccc | frame-relay-tcc | generic-services |
  multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc | vlan-vci-ccc | vlan-vpls);
```

Syntax for Logical Interfaces: SRX Series

```
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls
  | extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge |
  extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | frame-relay-port-ccc |
  vlan-ccc | vlan-vpls);
```

Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

[edit interfaces *interface-name*],
[edit interfaces rlsq *number*:*number*]

Logical Interfaces: SRX Series

[edit interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**flexible-ethernet-services**, **ethernet-ccc**, and **ethernet-tcc** options only).

Description

For M Series, MX Series, QFX Series, T Series, PTX Series, specify the physical link-layer encapsulation type.

For SRX Series, specify logical link layer encapsulation.

NOTE: Not all encapsulation types are supported on the switches. See the switch CLI.

Default

ppp—Use serial PPP encapsulation.

Physical Interface Options and Logical Interface Options

[Warning: element unresolved in stylesheets: <title> (in <config-options>). This is probably a new element that is not yet supported in the stylesheets.]

Physical Interface Options and Logical Interface Options

For physical interfaces:

NOTE: Frame Relay, ATM, PPP, SONET, and SATSOP options are not supported on EX Series switches.

- **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.
- **atm-pvc**—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).
- **cisco-hdlc**—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
  - CCC version (**cisco-hdlc-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
  - TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **cisco-hdlc-ccc**—Use Cisco-compatible HDLC framing on CCC circuits.
- **cisco-hdlc-tcc**—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.
- **ethernet-bridge**—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.
- **ethernet-over-atm**—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.
- **ethernet-tcc**—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.
- **ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.
- **ethernet-vpls-ppp**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.
- **ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.
- **extended-frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. When you use this encapsulation type, you can configure the **ccc** family only.
- **extended-frame-relay-ether-type-tcc**—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.
- **extended-frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.
- **extended-vlan-bridge**—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.
- **extended-vlan-ccc**—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.
- **extended-vlan-tcc**—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

- **extended-vlan-vpls**—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

  NOTE: **The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.**

- **flexible-ethernet-services**—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet interfaces, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.
- **flexible-frame-relay**—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- **frame-relay**—Use Frame Relay encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.
- **frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.

  NOTE: **When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.**

- **frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.

- **frame-relay-port-ccc**—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-tcc**—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **generic-services**—Use generic services encapsulation for services with a hierarchical scheduler.
- **multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.
-
- **ppp**—Use serial PPP encapsulation. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.
- **ppp-ccc**—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.
- **ppp-tcc**—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.
- **vlan-ccc**—Use Ethernet VLAN encapsulation on CCC circuits. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.

- **vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

NOTE:
- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
- Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure **family inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.

For logical interfaces:

- **frame-relay**—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- **multilink-frame-relay-uni-nni**—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- **ppp**—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- **ppp-over-ether**—This encapsulation is used for underlying interfaces of pp0 interfaces.

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

Related Documentation

- *Understanding Physical Encapsulation on an Interface*
- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- Configuring VLAN and Extended VLAN Encapsulation on page 256
- Configuring VLAN and Extended VLAN Encapsulation on page 256
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*

encryption (MACsec for MX Series)

Syntax encryption;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Enable MACsec encryption within a secure channel.

You can enable MACsec without enabling encryption. If a connectivity association with a secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic has not been tampered with and does not represent a security threat.

Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.

This command is used to enable encryption when MACsec is configured using secure association key (SAK) security mode only. When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the **no-encryption** configuration statement.

Default MACsec encryption is disabled when MACsec is configured using static SAK security mode, by default.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

enhanced-convergence

Syntax enhanced-convergence;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options mc-ae]
[edit interfaces irb unit *unit-number*]

Release Information Statement introduced in Junos OS Release 15.1R1.
Statement introduced in Junos OS Release 15.1X53-D60 for the QFX Series.

Description

NOTE: On EX9200 and QFX10000 switches, enhanced convergence is applicable for unicast traffic only—for example, when a MAC address is learned over an MC-AE interface, or an ARP entry is resolved over an MC-AE interface.

Improves Layer 2 and Layer 3 convergence time when a multichassis aggregated Ethernet (MC-AE) link goes down or comes up in a bridge domain or VLAN. Convergence time is improved because the traffic on the MC-AE interface is switched to the interchassis link (ICL) without waiting for a MAC address update.

If you have configured an IRB interface over an MC-AE interface that has enhanced convergences enabled, then you must configure enhanced convergence on the IRB interface as well. Enhanced convergence must be enabled for both Layer 2 and Layer 3 interfaces.

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

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Multichassis Link Aggregation*

ether-options

Syntax

```
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      port-priority
    }
  }
  (auto-negotiation | no-auto-negotiation);
  ethernet-switch-profile {
    tag-protocol-id;
  }
  (flow-control | no-flow-control);
  ieee-802-3az-eee;
  link-mode mode;
  (loopback | no-loopback);
  speed (speed | auto-negotiation);
}
```

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces interface-range *range*]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3R2.

Description

Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

ethernet (Chassis)

Syntax

```
ethernet {
  device-count number;
  lacp {
    link-protection {
      non-revertive;
    }
    system-priority;
  }
}
```

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure properties for Ethernet aggregated devices on the router.

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

Related Documentation

- Configuring Junos OS for Supporting Aggregated Devices on page 129
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

ethernet (Protocols OAM)

List of Syntax

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches on page 1146
Syntax: EX Series Switches and NFX Series Devices on page 1149

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches

```
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;
          }
        }
      }
    }
```


Syntax: EX Series Switches and NFX Series Devices

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      action {
        interface-down;
      }
      default-actions {
        interface-down;
      }
      event {
        adjacency-loss;
      }
    }
    esp-traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag (all |error | esp | interface | krt | lib |normal |task |timer);
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      mip-half-function (none | default |explicit);
      name-format (character-string | none | dns | mac+2oct);
      maintenance-association ma-name {
        continuity-check {
          hold-interval minutes;
          interface-status-tlv;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
          port-status-tlv;
        }
        mep mep-id {
          auto-discovery;
          direction down;
          interface interface-name;
          priority
          remote-mep mep-id {
            action-profile profile-name;
            sla-iterator-profile profile-name {
              data-tlv-size size;
              iteration-count count-value;
              priority priority-value;
            }
          }
        }
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
      }
    }
    performance-monitoring {
      sla-iterator-profiles {
        profile-name {
          calculation-weight {
            delay delay-value;
```

```
                delay-variation delay-variation-value;
              }
              cycle-time cycle-time-value;
              iteration-period iteration-period-value;
              measurement-type two-way-delay;
              passive;
            }
          }
        }
        traceoptions {
          file filename <files number> <match regex> <size size> <world-readable |
            no-world-readable>;
          flag flag ;
          no-remote-trace;
        }
      }
      link-fault-management {
        action-profile profile-name;
          action {
            syslog;
            link-down;
          }
          event {
            link-adjacency-loss;
            link-event-rate {
              frame-error  count;
              frame-period count;
              frame-period-summary  count;
              symbol-period count;
            }
          }
        interface interface-name {
          link-discovery (active | passive);
          pdu-interval interval;
          pdu-threshold threshold-value;
          remote-loopback;
          event-thresholds  {
            frame-error count;
            frame-period count;
            frame-period-summary  count;
            symbol-period count;
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
        }
        traceoptions {
          file filename <files number> <match regex> <size size> <world-readable |
            no-world-readable>;
          flag flag ;
          no-remote-trace;
        }
      }
    }
```

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in Junos OS Release 8.2 for MX, T, ACX Series routers, SRX firewalls, M320 and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.
connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 691
- *Example: Configuring Ethernet OAM Link Fault Management*

ethernet-policer-profile

Syntax

```
ethernet-policer-profile {
  input-priority-map {
    ieee802.1p premium [ values ];
  }
  output-priority-map {
    classifier {
      premium {
        forwarding-class class-name {
          loss-priority (high | low);
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

NOTE: On QFX Series standalone switches, this statement hierarchy is only supported on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure a class of service (CoS)-based policer. Policing applies to the inner VLAN identifiers, not to the outer tag. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the **premium** policer is not supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539

ethernet-ring

Syntax

```
ethernet-ring ring-name {
  control-vlan (vlan-id | vlan-name);
  data-channel {
    vlan number
  }
  east-interface {
    control-channel channel-name {
      vlan number;
      interface name interface-name
    }
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
  west-interface {
    control-channel channel-name {
      vlan number;
    }
  }
}
```

Hierarchy Level [edit protocols protection-group]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

Options *ring-name*—Name of the Ethernet protection ring.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

ethernet-switch-profile

Syntax

```
ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [values];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
    policer cos-policer-name {
      aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
    }
    storm-control storm-control-profile;
    tag-protocol-id tpid;
  }
  mac-learn-enable;
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 13.2X50-D15 for the EX Series switches.

Description

NOTE: On QFX Series standalone switches, the **ethernet-policer-profile** CLI hierarchy and the **mac-learn-enable** statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit

Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately. See CLI Explorer.

NOTE: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*]** hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

Default If the **ethernet-switch-profile** statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- Configuring MAC Address Filtering on page 544
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

evcs

Syntax

```
evcs evc-id {
  evc-protocol cfm;
  remote-uni-count count;
  multipoint-to-multipoint;
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 9.5.

Description On MX Series routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet virtual connection.

Options

remote-uni-count ***count***—(Optional) Specify the number of remote UNIs in the EVC configuration, the default is 1.

multipoint-to-multipoint—(Optional) Specify multiple points in the EVC configuration, the default is point-to-point if **remote-uni-count** is 1.

Remaining options 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 Ethernet Local Management Interface on page 627
- **lmi (Ethernet OAM) on page 1234**

evc-protocol cfm

Syntax

```
evc-protocol cfm {
  maintenance-association association-name | vpls routing-instance routing-id;
  maintenance-domain domain-id;
  mep-id mep-id;
}
```

Hierarchy Level [edit protocols oam ethernet evcs]

Release Information Statement introduced in Junos OS Release 9.5.
mep-id ***mep-id*** statement introduced in Junos OS Release 15.1.

Description Specify connectivity fault management (CFM) or virtual private LAN service (VPLS) as the Ethernet virtual connection (EVC) protocol.

Options

management-domain ***domain-id***—(Optional) For CFM, specify the CFM management domain.

management-association ***association-id***—(Optional) For CFM, specify the CFM management association.

routing-instance ***instance-id***—(Optional) For VPLS, specify the VPLS routing instance.

mep-id ***mep-id***—(Required for CFM) Identifier for the maintenance association endpoint

NOTE: This option is available on MX Series routers only.

Range: 1 through 8191

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

Related Documentation

- Configuring Ethernet Local Management Interface on page 627
- lmi (Ethernet OAM) on page 1234

event (LFM)

List of Syntax

Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches on page 1158
Syntax: EX Series Switches and NFX Series Devices on page 1158

Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches

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

Syntax: EX Series Switches and NFX Series Devices

```
event {
  link-adjacency-loss;
  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]

Release Information

Statement introduced in Junos OS Release 8.5 for MX, M, T, ACX Series routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX devices.

Description

Configure link events in an action profile for Ethernet OAM link fault management (LFM).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Monitoring Protocol Status on page 702
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

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. See CLI Explorer.

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 694

exclude-protocol (MX Series)

Syntax exclude-protocol *protocol-name*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a link using static connectivity association key (CAK) security mode.

When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.

Default Disabled.

All packets are secured on a link when MACsec is enabled, with the exception of all types of Spanning Tree Protocol (STP) packets.

Options ***protocol-name***—Specifies the name of the protocol that should not be MACsec-secured. Options include:

- **cdp**—Cisco Discovery Protocol.
- **lacp**—Link Aggregation Control Protocol.
- **lldp**—Link Level Discovery Protocol.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

exercise

Syntax request protection-group ethernet-aps exercise md <md> ma <ma>

Hierarchy Level [edit protocols protection-group ethernet-aps]

Description This configuration statement is used to test if APS is operating correctly, it does not interrupt regular APS operations.

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

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 213

failover-delay

Syntax failover-delay *milliseconds*;

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the failover delay for VRRP and VRRP for IPv6 operations.

Options ***milliseconds***—Specify the failover delay time, in milliseconds.

Range: 50 through 2000

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

Related Documentation

- Configuring VRRP and VRRP for IPv6 on page 346

family

Syntax

```
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      (input | output | input output);
    }
  }
  access-concentrator name;
  address address {
    ... the address subhierarchy appears after the main [edit interfaces interface-name unit
      logical-unit-number family family-name] hierarchy ...
  }
  bundle interface-name;
  core-facing;
  demux-destination {
    destination-prefix;
  }
  demux-source {
    source-prefix;
  }
  direct-connect;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
  }
  interface-mode (access | trunk);
  ipsec-sa sa-name;
  keep-address-and-control;
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  mtu bytes;
  multicast-only;
  negotiate-address;
  no-redirects;
  policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
  }
  primary;
  protocols [inet iso mpls];
  proxy inet-address address;
  receive-options-packets;
  receive-ttl-exceeded;
  remote (inet-address address | mac-address address);
  rpf-check {
```



```
    fail-filter filter-name
    mode loose;
  }
  sampling {
    input;
    output;
  }
  service {
    input {
      post-service-filter filter-name;
      service-set service-set-name <service-filter filter-name>;
    }
    output {
      service-set service-set-name <service-filter filter-name>;
    }
  }
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds> <filter [aci]>;
  (translate-discard-eligible | no-translate-discard-eligible);
  (translate-fecn-and-becn | no-translate-fecn-and-becn);
  translate-plp-control-word-de;
  unnumbered-address interface-name destination address destination-profile profile-name;
  vlan-id number;
  vlan-id-list [number number-number];
  address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    master-only;
    multipoint-destination address dlci dlci-identifier;
    multipoint-destination address {
      epd-threshold cells;
      inverse-arp;
      oam-liveness {
        up-count cells;
        down-count cells;
      }
      oam-period (disable | seconds);
      shaping {
        (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
          sustained rate);
        queue-length number;
      }
      vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    vrrp-groupgroup-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
```

```
                (preempt | no-preempt) {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        bandwidth-threshold bits-per-second priority-cost priority;
                        priority-cost priority;
                    }
                    priority-hold-time seconds;
                    route prefix routing-instance instance-name priority-cost priority;
                    }
                }
                virtual-address [ addresses ];
                }
            virtual-link-local-address ipv6-address;
        }
    }
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Option **max-sessions-vsa-ignore** introduced in Junos OS Release 11.4.

Description Configure protocol family information for the logical interface.

NOTE: Not all subordinate statements are available to every protocol family.

Options *family*—Protocol family:

- **any**—Protocol-independent family used for Layer 2 packet filtering

  NOTE: This option is not supported on T4000 Type 5 FPCs.

- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.
- **ethernet-switching**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ccc**—Circuit cross-connect protocol suite. You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—Internet Protocol version 4 suite. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—Internet Protocol version 6 suite. You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP).
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite. You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16. You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.
- **multilink-ppp**—Multilink Point-to-Point Protocol. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite. You can configure this protocol family for the logical interface of TCC physical interfaces.

- **tnp**—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only, as discussed in *Understanding Internal Ethernet Interfaces*.
- **vpls**—(M Series and T Series routers only) Virtual private LAN service. You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

  MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

  For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices*.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation
- *Configuring the Protocol Family*

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, EX Series switches, and MX series routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only using container interfaces) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.

NOTE:

- The fast APS switching feature is supported only within a single chassis on a MX series router using a container interface.
- 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*

fastether-options

Syntax

```
fastether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  ingress-rate-limit rate;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Fast Ethernet-specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Ethernet Interfaces Overview on page 3

flexible-vlan-tagging

Syntax flexible-vlan-tagging;

Hierarchy Level
[edit interfaces ae*x*],
[edit interfaces ge-*fpc/pic/port*],
[edit interfaces et-*fpc/pic/port*],
[edit interfaces ps0],
[edit interfaces xe-*fpc/pic/port*]

Release Information
Statement introduced in Junos OS Release 8.1.
Support for aggregated Ethernet added in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.

This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP.

This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX 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 VLAN Tagging on page 247
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 249
- *Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces*

flow-control

Syntax (flow-control | no-flow-control);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* multiservice-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Default Flow control is enabled.

NOTE: Flow control is enabled by default only on physical interfaces and it is disabled by default on aggregated Ethernet interfaces.

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

Related Documentation

- Configuring Flow Control on page 12
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

fnp

Syntax

```
fnp {
  interval <100ms | 1s | 10s | 1m | 10m>;;
  loss-threshold number
  interface interface name {
    domain-id domain-id
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Command introduced in Junos OS Release 11.4.

Description On routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet failure notification protocol.

Options

interval ***number***—Specifies the time between the transmission of FNP messages.

loss-threshold ***number***—FNP messages that can be lost before the FNP message is considered aged out and flushed.

interface ***interface-name***—Name of the Ethernet interface.

domain-id ***number***—Domain ID of the access network.

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

Related Documentation

- Ethernet Failure Notification Protocol Overview on page 734
- Configuring the Failure Notification Protocol on page 804

force switch

Syntax request protection-group ethernet-aps force-switch md <md> ma <ma>

Hierarchy Level [edit protocols protection-group ethernet-aps]

Description Forces traffic to switch from the active path to the alternate path. If the working path is the active path, traffic will be switched to the protection path. If the protection path is the active path, traffic will be switched to the protection path.

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

Related Documentation
- Ethernet Automatic Protection Switching Overview on page 213

force-up

Syntax force-up;

Hierarchy Level [edit interfaces *interface-name* aggregated-ethernet-options lacp]
[edit interfaces *interface-name* ether-options 802.3ad lacp]

Release Information Statement introduced in Junos OS Release 14.2 for MX Series routers.

Description Configure the peer interface (in MC-LAG) to remain up even with limited LACP capability.

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

Related Documentation
- *Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up*

forwarding-class (Gigabit Ethernet IQ Classifier)

Syntax

```
forwarding-class class-name {
  loss-priority (high | low);
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ interfaces only, define forwarding class name and option values.

Options

class-name—Name of forwarding class.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Specifying an Output Priority Map on page 541
- **input-priority-map on page 1200**
- **forwarding-class** statement in the *Class of Service Feature Guide for Routing Devices and EX9200 Switches*

forwarding-mode (100-Gigabit Ethernet)

Syntax

```
forwarding-mode {
  (sa-multicast | ...the following vlan-steering statement...);
  vlan-steering {
    vlan-rule (high-low | odd-even);
  }
}
```

Hierarchy Level [edit chassis fpc *slot* pic *slot*]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1 for MX Series routers.

Description Configure the interoperation mode for 100-Gigabit Ethernet PIC or the 100-Gigabit Ethernet MIC.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 448
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode*
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 457
- **sa-multicast (100-Gigabit Ethernet) on page 1350**
- **vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1443**
- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1444**

forwarding-mode (PTX Series Packet Transport Routers)

Syntax

```
forwarding-mode {
  sa-multicast
}
```

Hierarchy Level [edit chassis fpc *slot* pic *slot* port *port-number*]

Release Information Statement introduced in Junos OS Release 12.1X48R4.

Description Configure interoperability between 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP.

The remaining statement is explained separately. See CLI Explorer.

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

Related Documentation

- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 460
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 457

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 694
- Configuring Threshold Values for Fault Events in an Action Profile on page 703

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 694
- Configuring Threshold Values for Fault Events in an Action Profile on page 703

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 694
- Configuring Threshold Values for Fault Events in an Action Profile on page 703

framing (10-Gigabit Ethernet Interfaces)

Syntax framing (lan-phy | wan-phy);

Hierarchy Level [edit interfaces xe-*fpc/pic/port*]

[edit interfaces et-*fpc/pic/port*] (PTX Series Packet Transport Routers and MX Series Routers)

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 12.3R2 for PTX Series Packet Transport Routers.

Description For routers supporting the 10-Gigabit Ethernet interface, configure the framing format. WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600,T4000, and PTX Series Packet Transport Routers routers only.

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- On PTX Series routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ .
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.

Default Operates in LAN PHY mode.

Options **lan-phy**—10GBASE-R interface framing format that bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface.

wan-phy—10GBASE-W interface framing format that allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and SONET devices.

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

Related Documentation

- 10-Gigabit Ethernet Framing Overview on page 427
- *Configuring SONET Options for 10-Gigabit Ethernet Interfaces*

gigether-options

Syntax

```
gigether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (asynchronous-notification | no-asynchronous-notification);
  (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  fec
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  no-auto-mdix
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
  speed
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [ tpids ];
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
```

```
    }
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Gigabit Ethernet specific interface properties.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- Ethernet Interfaces Overview on page 3
- *gigether-options (ACX Series)*

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

Hierarchy Level [edit interfaces *interface-name*]
[edit interfaces **interface-range** *interface-range-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default Updating of the ARP cache is disabled on all Ethernet interfaces.

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

Related Documentation
- Configuring Gratuitous ARP on page 19
- **no-gratuitous-arp-request on page 1274**

guest-vlan (MX Series in Enhanced LAN Mode)

Syntax guest-vlan (*vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Specify the VLAN to which an interface is moved when no 802.1X supplicants are connected on the interface. The VLAN specified must already exist on the switch.

Default None

Options ***vlan-id***—VLAN tag identifier of the guest VLAN.

vlan-name—Name of the guest VLAN.

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

guard-interval

Syntax guard-interval *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL link receives notification, restores the link, and waits for the restore interval before issuing another block on the same link. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Options *number*—Guard timer interval, in milliseconds.

Range: 10 through 2000 ms

Default: 500 ms

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

hold-interval (Protection Group)

Syntax hold-interval *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify the hold-off timer interval *for all rings* in 100 millisecond (ms) increments.

Options *number*—Hold-timer interval, in milliseconds.

Range: 0 through 10,000 ms

Default: 100 ms

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

hold-multiplier

Syntax hold-multiplier *number*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a value for the LLDP hold multiplier.

Hold timer interval in seconds to cache learned LLDP information before discarding.

Options *number*—Advertisement interval multiplier for LLDP cache discard.

Default: 4 (giving 120 second LLDP cache lifetime with other defaults)

Range: 2 through 10

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

Related Documentation

- Configuring LLDP on page 338

hold-time up

Syntax hold-time up *timer-value*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options lacp],

Release Information Statement introduced in Junos OS Release 14.2R3.

Description Specifies the time period for which the Link Aggregation Control Protocol (LACP) maintains the state of a child (member) link as expired or default.

When a child link goes from the current state to the expired state, the LACP monitors the reception of protocol data units (PDUs) on the child link for the configured hold-up time interval and does not allow the child link to transition back to the current state. This configuration thus prevents excessive flapping of a child link on an aggregated Ethernet interface.

The configured hold-up timer value is applicable to all the child links within a link aggregated (LAG) interface. By default, this feature is disabled.

Options ***timer-value***—Hold-up interval in seconds.

Range: 1 to 6000 seconds

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

Related Documentation

- Configuring LACP for Aggregated Ethernet Interfaces on page 140
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*

iccp

Syntax

```
iccp {
  traceoptions; {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
      <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
  local-ip-address ip address;
  session-establishment-hold-time value;
  authentication-key string;
  peer ip-address {
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    redundancy-group-id-list redundancy-group-id-list;
    liveness-detection;
  }
}
```

Hierarchy Level

[edit protocols iccp]
[edit logical-systems *logical-system-name* protocols iccp]

Release Information

Statement introduced in Junos OS Release 10.0.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

Default

If you do not include this statement, no ICCP protocol tracing operations are performed.

Options

traceoptions—Set Interchassis Control Protocol (ICCP) tracing options.

local-ip-address—Specify the source address where the ICCP packet is routed.

session-establishment-hold-time—Specify if the chassis takes over as the master at the ICCP session.

authentication-key —Specify TCP Message Digest 5 (MD5) option for an ICCP TCP session.

peer ip-address—Specify the IP address of the peer that hosts an MC-LAG. You must configure ICCP for both peers that host the MC-LAG.

redundancy-group-id-list—Specify the redundancy groups between two ICCP peers.

liveness-detection—Specify Bidirectional Forwarding Detection (BFD) protocol options.

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

Related Documentation
- *Configuring ICCP for MC-LAG*

id (MACsec for MX Series)

Syntax
id {
  mac-address *mac-address*;
  port-id *port-id-number*;
}

Hierarchy Level
[edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name*]

Release Information
Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description
Specify a MAC address and a port that traffic on the link must be from to be accepted by the interface when MACsec is enabled using static secure association key (SAK) security mode.

Options
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 Media Access Control Security (MACsec) on MX Series Routers*

ieee802.1p

Syntax

ieee802.1p premium [*values*];

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]

Release Information

Statement introduced before Junos Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, configure premium priority values for IEEE 802.1p input traffic.

Options

values—Define IEEE 802.1p priority values to be treated as premium.

Range: 0 through 7

Required Privilege Level

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

Related Documentation

- Specifying an Input Priority Map on page 541

igmp-snooping

List of Syntax

Syntax (EX Series, QFX Series, and NFX Series) on page 1189
Syntax (MX Series) on page 1189
Syntax (SRX Series) on page 1191

Syntax (EX Series, QFX Series, and NFX Series)

```
igmp-snooping {
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag (detail | disable | receive | send);
  }
  vlan (vlan-name | all) {
    data-forwarding {
      source {
        groups group-prefix;
      }
      receiver {
        source-vlans vlan-list;
        install ;
      }
    }
    disable;
    immediate-leave;
    interface interface-name {
      group-limit limit;
      host-only-interface;
      immediate-leave;
      multicast-router-interface;
      static {
        group multicast-ip-address;
      }
    }
    l2-querier {
      source-address ip-address;
      }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    version number;
  }
}
```

Syntax (MX Series)

```
igmp-snooping {
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    immediate-leave;
```

```
      multicast-router-interface;
      static {
        group ip-address {
          source ip-address;
        }
      }
    }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
    vlan vlan-id {
      immediate-leave;
      interface interface-name {
        group-limit limit;
        host-only-interface;
        immediate-leave;
        multicast-router-interface;
        static {
          group ip-address {
            source ip-address;
          }
        }
      }
      proxy {
        source-address ip-address;
      }
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
```


Syntax (SRX Series)

```
igmp-snooping {
  vlan (all | vlan-name) {
    immediate-leave;
    interface interface-name {
      group-limit range;
      host-only-interface;
      multicast-router-interface;
      immediate-leave;
      static {
        group multicast-ip-address {
        source ip-address;
        }
      }
    }
    l2-querier {
      source-address ip-address;
      }
    proxy {
      source-address ip-address;
    }
    qualified-vlan vlan-id;
    query-interval number;
    query-last-member-interval number;
    query-response-interval number;
  robust-count number;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
    }
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]
[edit routing-instances *routing-instance-name* protocols]
[edit protocols]

Release Information

Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 18.1R1 for SRX1500 devices.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

Configure IGMP snooping to constrain multicast traffic to only the ports that have receivers attached. IGMP snooping enables the device to selectively send out multicast packets on only the ports that need them. Without IGMP snooping, the device floods the packets on every port. The device listens for the exchange of IGMP messages by the device and the end hosts. In this way, the device builds an IGMP snooping table that has a list of all the ports that have requested a particular multicast group. The factory default configuration enables IGMP snooping on all VLANs.

NOTE: IGMP snooping must be disabled on the device before enabling ISSU.

Default IGMP snooping is disabled on the device.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *IGMP Snooping in MC-LAG Active-Active Mode*
- *Example: Configuring IGMP Snooping on SRX Series Devices*
- *IGMP Snooping Overview*

ignore-l3-incompletes

Syntax ignore-l3-incompletes;

Hierarchy Level [edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description Ignore the counting of Layer 3 incomplete errors on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.

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

Related Documentation

- Ignoring Layer 3 Incomplete Errors on page 19

include-sci (MACsec for MX Series)

Syntax include-sci;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies that the SCI tag should be appended to each packet on a link that has enabled MACsec.

SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an MX240, MX480, or MX960 router. This option is, therefore, redundant to be configured.

This option is used only when connecting a router to a host device that requires SCI tagging. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

Default SCI tagging is enabled on MX Series routers that have enabled MACsec using static connectivity association key (CAK) security mode, by default.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

ingress-policer-overhead

Syntax ingress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1.
Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

Description Add the configured number of bytes to the length of a packet entering the interface.

Configure a policer overhead to control the rate of traffic received on an interface. Use this feature to help prevent denial-of-service (DoS) attacks or to enforce traffic rates to conform to the service-level agreement (SLA). When you configure a policer overhead, the configured policer overhead value (bytes) is added to the length of the final Ethernet frame. This calculated length of frame is used to determine the policer or the rate-limiting action.

Traffic policing combines the configured policy bandwidth limits and the burst size to determine how to meter the incoming traffic. If you configure a policer overhead on an interface, Junos OS adds those bytes to the length of incoming Ethernet frames. This added overhead fills each frame closer to the burst size, allowing you to control the rate of traffic received on an interface.

You can configure the policer overhead to rate-limit queues and Layer 2 and Layer 3 policers, for standalone (SA) and high-avalability (HA) deployments. The policer overhead and the shaping overhead can be configured simultaneously on an interface.

NOTE: vSRX supports policer overhead on Layer 3 policers only.

The policer overhead applies to all interfaces on the PIC. In the following example, Junos OS adds 10 bytes of overhead to all incoming Ethernet frames on ports ge-0/0/0 through ge-0/0/4.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 10
```

NOTE: vSRX only supports fpc 0 pic 0. When you commit the ingress-policer-overhead statement, the vSRX takes the PIC offline and then back online.

You need to craft the policer overhead size to match your network traffic. A value that is too low will have minimal impact on traffic bursts. A value that is too high will rate-limit too much of your incoming traffic.

In this example, the policer overhead of 255 bytes is configured for ge-0/0/0 through ge-0/0/4. The firewall policer is configured to discard traffic when the burst size is over 1500 bytes. This policer is applied to ge-0/0/0 and ge 0/0/1. Junos OS adds 255 bytes to every Ethernet frame that comes into the configured ports. If, during a burst of traffic, the combined length of incoming frames and the overhead bytes exceeds 1500 bytes, the policer starts to discard further incoming traffic.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 255
set interfaces ge-0/0/0 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/0 unit 0 family inet address 10.9.1.2/24
set interfaces ge-0/0/1 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/1 unit 0 family inet address 10.9.2.2/24
set firewall policer overhead_policer if-exceeding bandwidth-limit 32k
set firewall policer overhead_policer if-exceeding burst-size-limit 1500
set firewall policer overhead_policer then discard
```

Options ***bytes***—Number of bytes added to a frame entering an interface.

Range: 0–255 bytes

Default: 0

```
[edit chassis fpc 0 pic 0]
user@host# set ingress-policer-overhead 10;
```

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

Related Documentation

- *ingress-shaping-overhead*
- *Policer Overhead to Account for Rate Shaping Overview*
- *Example: Configuring Policer Overhead to Account for Rate Shaping*
- *Configuring a Policer Overhead*
- *CoS on Enhanced IQ2 PICs Overview*

ingress-rate-limit

Syntax ingress-rate-limit *rate*;

Hierarchy Level [edit interfaces *interface-name* fastether-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Perform port-based rate limiting on ingress traffic arriving on Fast Ethernet 8-port, 12-port, and 48-port PICs.

Options *rate*—Traffic rate, in megabits per second (Mbps).

Range: 1 through 100 Mbps

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

Related Documentation

- Configuring the Ingress Rate Limit on page 10

inner-tag-protocol-id

Syntax inner-tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options** ethernet-switch-profile tag-protocol-id **[*tpids*]]** hierarchy level.

On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.

Default If the **inner-tag-protocol-id** statement is not configured, the TPID value is 0x8100.

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

Related Documentation

- Configuring Inner and Outer TPIDs and VLAN IDs on page 566

inner-vlan-id

Syntax inner-vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers or 100-Gigabit Ethernet Type 5 PIC with CFP, or on Ethernet interfaces on EX Series switches, specify the VLAN ID to rewrite for the inner tag of the final packet.

You cannot include the **inner-vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement and the **inner-vlan-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **inner-vlan-id** statement you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Options *number*—VLAN ID number.

Range: 0 through 4094

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

Related Documentation

- Configuring Inner and Outer TPIDs and VLAN IDs on page 566

input-policer

Syntax input-policer *policer-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* layer2-policer]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Apply a single-rate two-color policer to the Layer 2 input traffic at the logical interface. The **input-policer** and **input-three-color** statements are mutually exclusive.

Options *policer-name*—Name of the single-rate two-color policer that you define at the **[edit firewall]** hierarchy level.

Usage Guidelines See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- Example: Configuring Gigabit Ethernet Policers on page 544
- **input-three-color on page 1201**
- **layer2-policer on page 1219**
- **logical-interface-policer on page 1238**
- **output-policer on page 1291**
- **output-three-color on page 1293**

input-priority-map

Syntax

```
input-priority-map {
  ieee802.1p premium [ values ];
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the input policer priority map to be applied to incoming frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Specifying an Input Priority Map on page 541
- output-priority-map on page 1292

input-three-color

Syntax input-three-color *policer-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* layer2-policer]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Apply a single-rate or two-rate three-color policer to the Layer 2 input traffic at the logical interface. The **input-three-color** and **input-policer** statements are mutually exclusive.

Options ***policer-name***—Name of the single-rate or two-rate three-color policer.

Usage Guidelines See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- Example: Configuring Gigabit Ethernet Policers on page 544
- **input-policer on page 1199**
- **layer2-policer on page 1219**
- **logical-interface-policer on page 1238**
- **output-policer on page 1291**
- **output-three-color on page 1293**

input-vlan-map (Aggregated Ethernet)

Syntax

```
input-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 8.2.
Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

Description

Define the rewrite profile to be applied to incoming frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet inferfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Stacking a VLAN Tag on page 569
- **output-vlan-map (Aggregated Ethernet) on page 1294**

input-vlan-map

Syntax

```
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.

The remaining statements are explained separately. See CLI Explorer.

NOTE: Connectivity fault management (CFM) sessions for all interfaces in which **input-vlan-map** is configured are supported only if the interface also has an explicit configuration for **output-vlan-map** as **output-vlan-map pop;**. See output-vlan-map. This configuration is required for all the interfaces in the topology even when the CFM session is on that interface or on a different interface in the data path of the same topology.

Required Privilege Level

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

Related Documentation

- Stacking a VLAN Tag on page 569
- **output-vlan-map on page 1295**
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

interface

Syntax

```
interface (all | interface-name) {
  disable;
}
```

Hierarchy Level

[edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Specify an LLDP interface.

Options

interface-name—A valid physical interface name.

NOTE: On MX Series and T Series routers, you run LLDP on a physical interface, such as ge-1/0/0, and not at the logical interface (unit) level.

Starting with Junos OS Release 14.2, you can also specify LLDP neighbor details for management interfaces, such as fxp or me, on MX Series routers.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

all—Run LLDP on all interfaces.

disable—Disable LLDP on the specified interface

Required Privilege Level

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

Related Documentation

- Configuring LLDP on page 338

interface (IEEE 802.1x)

Syntax

```
interface interface-id {
  maximum-requests integer;
  quiet-period seconds;
  reauthentication (disable | interval seconds);
  retries integer;
  server-timeout seconds;
  supplicant (single);
  supplicant-timeout seconds;
  transmit-period seconds;
}
```

Hierarchy Level [edit protocols dot1x authenticator]

Release Information Statement introduced in Junos OS Release 9.3.

Description Use this statement to configure the 802.1x Port-Based Network Access Control protocol-specific Ethernet interface options.

Default The default values are provided for the options below on the respective statement pages.

Options

maximum-requests—Specify the maximum number of retransmission times for an EAPOL Request packet to the client before it times out the authentication session.

quiet-period—Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting the authentication.

reauthentication—Includes two options:

- **disable**—Periodic reauthentication of the client is disabled.
- **interval**—Specify the periodic reauthentication time interval.

retries—Specify the number of tries after which the port remains in the wait state for **quiet-period** seconds before reattempting the authentication.

server-timeout—Specify the number of seconds the port waits for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

supplicant (*single*)—Specify supplicant single mode. See the usage guidelines to configure other modes.

supplicant-timeout—Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

transmit-period—Specify the number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Required Privilege Level

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authenticator on page 1087**
- **dot1x on page 1124**

interface (OAM Link-Fault Management)

List of Syntax

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 1207
Syntax: EX Series Switches and NFX Series Devices on page 1207

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches

```
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;
  }
}
```

Syntax: EX Series Switches and NFX Series Devices

```
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
```

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

Release Information Statement introduced in Junos OS Release 8.2 for T, M, MX and ACX Series Routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

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

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 691
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

interface (Static MAC Bypass)

Syntax interface [*interface-names*];

Hierarchy Level [edit protocols authentication-access-control]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure interfaces on which the specified MAC addresses are allowed to bypass RADIUS authentication and allowed to connect to the LAN without authentication.

Options ***interface-names***—List of interfaces.

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

Related Documentation

interfaces (MACsec for MX Series)

Syntax

```
interfaces interface-name {
   connectivity-association connectivity-association-name;
}
```

Hierarchy Level [edit security macsec]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Applies the specified connectivity association to the specified interface to enable MACsec.

One connectivity association can be applied to multiple interfaces.

You must always use this statement to apply a connectivity association to an interface to enable MACsec. You must complete this configuration step regardless of whether MACsec is enabled using static connectivity association key (CAK) security mode or static secure association key (SAK) security mode.

If you are enabling MACsec using static SAK security mode and need to configure MACsec on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is then applied to the interface using this statement to enable MACsec for traffic entering and leaving the interface.

NOTE: Starting in Junos OS Release 16.1R2, when Media Access Control Security (MACsec) is enabled on an interface, the interface flow control capability is enabled by default, regardless of the configuration that you set using the **(flow-control | no-flow-control)** statement at the **[edit interfaces *interface- name* gigether-options]** hierarchy level. When MACsec is disabled, interface flow control is restored to the configuration that you set using the **flow-control** statement at the **[edit interfaces]** hierarchy level. When MACsec is enabled, additional header bytes are added to the packet by the MACsec PHY. With line rate traffic, when MACsec is enabled and flow control is disabled, the pause frames sent by the MACsec PHY are terminated by the MIC's MAC (enhanced 20-port Gigabit Ethernet MICs on MX Series routers) and not transferred to the Packet Forwarding Engine, causing framing errors. Therefore, when MACsec is enabled on an interface, flow control is also automatically enabled on such an interface.

Default Interfaces are not associated with any connectivity associations, by default.

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

interface-group

Syntax

```
interface-group {
  interface-device-name
  unit-list
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 18.1R1.

Description

Mark the interface group down for the action profile configured with the action **interface-group-down**. Provides information for the interface-group on which the configured action will be taken when the configured event occur for a specific remote MEP ID.

Options

interface-device name—Name of the interface device. Only Ethernet devices are allowed. The device interface name includes **ge, ae, xe and et.** .

unit-list—One or more logical interface unit numbers.

Range: A string in the range <0-16385> or <0-16385>-<0-16385>. For example, **unit-list**[12 23-33 44]

Required Privilege Level

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

Related Documentation

- Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces on page 888
- **interface-group-down on page 1211**

interface-group-down

Syntax interface-group-down

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *action-profile-name* action]

Release Information Statement introduced in Junos OS Release 18.1R1.

Description Mark the interface group down.

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

Related Documentation

- Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces on page 888
- interface-group on page 1210

interface-none

Syntax interface-none;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name* east-interface]

[edit protocols protection-group ethernet-ring *ring-name* west-interface]

Description Designates port as not used for Ethernet ring protection.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

isolated-vlan (MX Series)

Syntax isolated-vlan *vlan-id*;

Hierarchy Level [edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* ,
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers.

Description Configure the specified isolated VLAN to be a secondary VLAN of the specified primary VLAN. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.

NOTE: When you specify this configuration statement, the VLAN ID of a logical interface that you associate with a bridge domain that matches with the VLAN ID that you specify using the **isolated-vlan** state is treated as an isolated port.

Options *vlan-id*—Individual VLAN IDs separated by a space.

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

Related Documentation

key (MACsec for MX Series)

Syntax key *key-string*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name* security-association *security-association-number*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies the static security key to exchange to enable MACsec using static secure association key (SAK) security mode.

The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec when enabling MACsec using SAK security mode.

You must configure at least two security associations with unique security association numbers and key strings to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.

Default This statement does not have a default value.

Options ***key-string***—Specifies the key to exchange with the other end of the link on the secure channel. The *key-string* is a 32-digit hexadecimal string that is created by the user.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

key-server-priority (MACsec for MX Series)

Syntax key-server-priority *priority-number*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* mka]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.

The switch with the lower *priority-number* is selected as the key server.

If the *priority-number* is identical on both sides of a point-to-point link, the MKA protocol selects the device with the lower MAC address as the key server.

Default The default key server priority number is 16.

Options ***priority-number***—Specifies the MKA server election priority number.

The *priority-number* can be any number between 0 and 255. The lower the number, the higher the priority.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

lacp (802.3ad)

Syntax

```
lacp {
  port-priority port-priority;
}
```

Hierarchy Level

[edit interfaces *interface-name* fastether-options 802.3ad],
[edit interfaces *interface-name* gigether-options 802.3ad]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure the Link Aggregation Control Protocol (LACP) port priority for Ethernet interfaces.

Options

port-priority—Priority for being elected as the active port to collect and distribute traffic. A smaller value indicates a higher priority for selection.

Range: 0 through 65,535

Default: 127

Required Privilege Level

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

Related Documentation

- Configuring LACP for Aggregated Ethernet Interfaces on page 140
- **port-priority on page 1311**

lacp (Aggregated Ethernet)

List of Syntax

Syntax (NFX Series) on page 1216
Syntax (EX Series) on page 1216

Syntax (NFX Series)

```
lacp (active | passive) {
  admin-key key;
  fast-failover;
  link-protection {
    disable;
    (revertive | non-revertive);
  }
  periodic interval
  system-ID mac-address;
  system-priority priority;
  force-up;
}
```

Syntax (EX Series)

```
lacp {
  (active | passive);
  admin-key key;
  accept-data;
  fast-failover;
    link-protection {
    disable;
    (revertive |non-revertive);
  }
  periodic interval;
  system-id mac-address;
  system-priority priority;
}
```

Hierarchy Level (EX Series)

[edit interfaces ae*x* aggregated-ether-options]
[edit logical-systems *logical-system-name* interfaces ae*X* aggregated-ether-options]

Hierarchy Level (NFX Series)

[edit interfaces *interface-name*aggregated-ether-options]

Release Information

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

Description

Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.

For EX Series, when you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

NOTE: The force-up statement is not supported on QFX10002 switches.

Default If you do not specify LACP as either **active** or **passive**, LACP remains passive.

Options **active**—Initiate transmission of LACP packets.

admin-key ***number***—Specify an administrative key for the router or switch.

NOTE: You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive—Respond to LACP packets.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- *Configuring Link Aggregation*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- Configuring LACP for Aggregated Ethernet Interfaces on page 140

layer2-policer

Syntax

```
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series, MX Series, and T Series routers, and for aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on EX Series switches, apply Layer 2 logical interface policers. The following policers are supported:

- Two-color
- Single-rate tricolor marking (srTCM)
- Two-rate tricolor marking (trTCM)

Two-color and tricolor policers are configured at the **[edit firewall]** hierarchy level.

Options

input-policer ***policer-name***—Two-color input policer to associate with the interface. This statement is mutually exclusive with the **input-three-color** statement.

input-three-color ***policer-name***—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the **input-policer** statement.

output-policer ***policer-name***—Two-color output policer to associate with the interface. This statement is mutually exclusive with the **output-three-color** statement.

output-three-color ***policer-name***—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the **output-policer** statement.

Required Privilege Level

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

Related Documentation

- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 546

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 701

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 692

link-degrade-monitor

Syntax

```
link-degrade-monitor {
  actions media-based;
  recovery {
    (auto | manual);
    timer timer;
  }
  thresholds {
    clear clear-value;
    interval interval-value;
    set set-value;
    warning-clear warning-clear-value;
    warning-set warning-set-value;
  }
}
```

Hierarchy Level [edit interfaces *interfaces- name*]

Release Information Statement introduced in Junos OS Release 15.1.

Description Configure link degrade monitoring on an interface and specify the corrective action to be triggered when a link degrade event is detected. Deleting the configuration disables the feature. When configured, the feature monitors the quality of physical links on Ethernet interfaces (10-Gigabit, 40-Gigabit, and 100-Gigabit) and triggers the user-configured action when the link's bit error rate (BER) value breaches the preconfigured threshold. This feature can detect a BER value as low as 10^{-13} to 10^{-5}.

Options **actions media based**—Action to be taken when a link degrade event is detected. A media-based action brings down the physical link at both local and remote ends of the interface, and stops BER monitoring at the local end until an autorecovery is triggered.

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

- Link Degrade Monitoring Overview on page 381
- thresholds on page 1395
- recovery on page 1339
- request interface link-degrade-recover on page 1495

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 700

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 703

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. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 691

link-mode

Syntax link-mode *mode* (automatic | full-duplex | half-duplex);

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* ether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description Set the device's link connection characteristic.

Options *mode*—Link characteristics:

- **automatic**—Link mode is negotiated. This is the default for EX Series switches.
- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.

Default: Fast Ethernet interfaces can operate in either full-duplex or half-duplex mode. The router's or switch's management Ethernet interface, **fxp0** or **em0**, and the built-in Fast Ethernet interfaces on the FIC (M7i router) autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.

NOTE: On EX Series switches, if **no-auto-negotiation** is specified in **[edit interfaces *interface-name* ether-options]**, you can select only **full-duplex** or **half-duplex**. If **auto-negotiation** is specified, you can select any mode.

NOTE:

- Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.
- Starting with Junos OS release 17.4R1 and later, the **link-mode** configuration is not supported for 10-Gigabit Ethernet interfaces.

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

Related Documentation

- Configuring the Link Characteristics on Ethernet Interfaces on page 13
- *Understanding Management Ethernet Interfaces*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*

link-protection

Syntax

```
link–protection {
  disable;
  (revertive |non-revertive);
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit interfaces ae*x* aggregated-ether-options *lacp*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.
Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the **[edit interfaces ge-*fpc/pic/port* gigether-options 802.3ad ae*x*]** hierarchy level or the **[edit interfaces fe-*fpc/pic/port* fastether-options 802.3ad ae*x*]** hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify **link-protection** at the following hierarchy levels:

- **[edit interfaces ge-*fpc/pic/port* ether-options 802.3ad ae*x*]**
- **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level or at the **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level.

  To disable link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the **[edit interfaces ae*x* aggregated-ether-options]** hierarchy level or the **[edit interfaces ae*x* aggregated-ether-options *lacp*]]** hierarchy level.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 148
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

link-protection (non-LACP)

Syntax

```
link–protection {
  link-protection-revertive;
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 17.3R1.

Description User can specify the *link-protection-revertive* statement in the link protection configuration at the aggregated Ethernet interface level to set revertive mode. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in recalculation of the priorities and traffic will switch from the currently active link to the newly added, higher-priority link. Recalculation of priorities is performed only while link event such as addition\deletion and UP/Down operation on link, that is, configuration of this option will not result in any recalculation immediately until next link-event occurs.

In addition to enabling static link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the **[edit interfaces ge-*fpc/pic/port* gigether-options 802.3ad ae*x*]** hierarchy level or the **[edit interfaces fe-*fpc/pic/port* fastether-options 802.3ad ae*x*]** hierarchy level.

For static link protection, specify **link-protection** at the following hierarchy levels:

- **[edit interfaces ge-*fpc/pic/port* ether-options 802.3ad ae*x*]**
- **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level or at the **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level.

  To disable static link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the **[edit interfaces ae*x* aggregated-ether-options]** hierarchy level.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 148

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level (EX Series) [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
mixed option added in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers and 15.1F6 and 16.1R2 for PTX3000 routers.

Description For aggregated Ethernet interfaces only, set the required link speed.

Options ***speed***—For aggregated Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Aggregated Ethernet links on T Series, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.

- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

mixed—Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface.

Required Privilege Level

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

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 102
- Configuring Aggregated Ethernet Link Speed on page 134
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 124
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

link-speed (Aggregated SONET/SDH)

Syntax link-speed (*speed* | mixed);

Hierarchy Level [edit interfaces as*x* aggregated-sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.
mixed option added in Release 8.0.

Description For aggregated SONET/SDH interfaces only, set the required link speed.

Options *speed*—Aggregated SONET/SDH links can have one of the following speed values.

- **oc3**—Links are OC3c or STM1c.
- **oc12**—Links are OC12c or STM4c.
- **oc48**—Links are OC48c or STM16c.
- **oc192**—Links are OC192c or STM64c.
- **oc768**—Links are OC768c or STM256c.

mixed—For aggregated SONET/SDH links on T Series routers, you can mix interface speeds in SONET/SDH aggregation bundles. Interface speeds from OC3 through OC768 are supported.

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

Related Documentation
- Configuring Aggregated Ethernet Link Speed on page 134
- *Configuring Aggregated SONET/SDH Interfaces*

lldp

Syntax

```
lldp {
  advertisement-interval seconds;
  disable;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  management-address ip-management-address;;
  mau-type
  port-description-type {
    interface-alias;
    interface-description;
  }
  port-id-subtype {
    interface-name;
    locally-assigned;
  }
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  transmit-delay seconds;
}
```

Hierarchy Level [edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify LLDP configuration parameters.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring LLDP on page 338

lldp-configuration-notification-interval

Syntax lldp-configuration-notification-interval *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a time for the period of SNMP trap notifications to the Master Agent to wait regarding changes in database information.

Options *seconds*—Time for the period of SNMP trap notifications about the LLDP database. This feature is disabled by default.

Range: 0 through 3600

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

Related Documentation

- Configuring LLDP on page 338

lmi (Ethernet OAM)

Syntax

```
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 ethernet]

Release Information Statement introduced in Junos OS Release 9.5.

Description On routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet Local Management Interface (E-LMI).

NOTE: On MX Series routers, E-LMI is supported on Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), and Aggregated Ethernet (ae) interfaces configured on MX Series routers with DPC only.

Options

status-counter *count*—Status counter (N393), defaults to 4.

interface *name*—Polling verification timer (T392), defaults to 15 seconds.

uni-id *uni-name*—(Optional) Defaults to the physical interface name.

status-counter *number*—(Optional) Defaults to a global value.

polling-verification-timer *value*—(Optional) Defaults to a global value.

evc-map-type (all-to-one-bundling | bundling | service-multiplexing)—Specify the Ethernet virtual connection (EVC) map type.

evc *evc-name*—Specify the name of the EVC.

default-evc—Set the specified EVC as the default EVC.

vlan-list *vlan-id-list*—Specify a group of VLANs to assign to the EVC.

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

Related Documentation

- Configuring Ethernet Local Management Interface on page 627
- evcs on page 1156

load-balance

Syntax

```
load-balance {
  adaptive{
    pps;
    scan-interval multiple;
    tolerance percentage;
  }
  no-adaptive;
  per-packet;
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 13.3.
Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.

Description Load-balances the received traffic across all the available paths of aggregated Ethernet bundles for better link utilization.

Options

adaptive— (MX Series and PTX Series) Corrects a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an Aggregated Ethernet bundle.

no-adaptive— (MX Series and PTX Series) Disables the adaptive load-balancing solution configured to distribute traffic by using a feedback mechanism.

per-packet— (MX Series only) Randomly sprays packets to the aggregate next hops in a round-robin manner to avoid traffic imbalance.

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

Related Documentation

- Understanding Aggregated Ethernet Load Balancing on page 158
- Example: Configuring Aggregated Ethernet Load Balancing on page 163

load-balance-stateful (Aggregated Ethernet Interfaces)

Syntax

```
load-balance-stateful {
  per-flow;
    rebalance interval;
  load-type (low | medium | large);
}
```

Hierarchy Level [edit interfaces ae*X* *unit logical-unit-number* forwarding-options]

Release Information Statement introduced in Junos OS Release 13.2R1.

Description Define the capability to perform uniform load balancing and also perform rebalancing is introduced on MX Series routers with MPCs, except MPC3Es and MPC4Es. Rebalancing is not supported when load-balancing is skewed or distorted owing to a change in the number of flows. The mechanism to record and retain states for the flows and distribute the traffic load accordingly is added. As a result, for m number of flows, they are distributed among n member links of a LAG bundle or among the unilist of next-hops in an ECMP link. This method of splitting the load among member links is called stateful load balancing and it uses 5-tuple information (source and destination addresses, protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

Options **stateful**—Define the stateful load-distribution mechanism for traffic flows on aggregated Ethernet interfaces.

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

Related Documentation

- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces on page 181

load-type (Aggregated Ethernet Interfaces)

Syntax load-type (low | medium | large);

Hierarchy Level [edit interfaces ae*X* *unit logical-unit-number* forwarding-options load-balance-stateful]

Release Information Statement introduced in Junos OS Release 13.2R1.

Description Define the load-balancing type to inform the Packet Forwarding Engine regarding the appropriate memory pattern to be used for traffic flows. The approximate number of flows for effective load-balancing for each keyword is a derivative.

Options **low**—Define a low load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is less or minimal (between 1 and 100 flows).

medium—Define a medium or moderate load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is relatively higher (between 100 and 1000 flows).

large—Define a high load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is excessive or reaches the maximum supported flows (between 1000 and 10,000 flows).

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

Related Documentation

- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces on page 181

lockout

Syntax request protection-group ethernet-aps lockout md <md> ma <ma>

Hierarchy Level [edit protocols protection-group ethernet-aps]

Description Configure a lockout of the protection path, forcing the use of the working path and locking out the protect path regardless of anything else.

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

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 213

logical-interface-policer

Syntax logical-interface-policer;

Hierarchy Level [edit dynamic-profiles *profile-name* firewall policer *policer-name*],
[edit dynamic-profiles *profile-name* firewall three-color-policer *name*],
[edit firewall atm-policer*atm-policer-name*],
[edit firewall policer *policer-name*],
[edit firewall policer *policer-template-name*],
[edit firewall three-color-policer *policer-name*],
[edit logical-systems *logical-system-name* firewall policer *policer-name*],
[edit logical-systems *logical-system-name* firewall three-color-policer *name*]

Release Information Statement introduced before Junos OS Release 7.4.
Support at the **[edit firewall three-color-policer *policer-name*]** hierarchy level introduced in Junos OS Release 8.2.
Logical systems support introduced in Junos OS Release 9.3.
Support at the **[edit dynamic-profiles ... policer *policer-name*]** and **[edit dynamic-profiles ... three-color-policer *name*]**hierarchy levels introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure a logical interface policer.

NOTE: Starting in Junos OS Release 12.2R2, on T Series Core Routers only, you can configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. You must include the **logical-interface-policer** statement to do so.

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

Related Documentation

- *Two-Color and Three-Color Logical Interface Policers*
- *Traffic Policer Types*
- *Configuring and Applying Tricolor Marking Policers*
- *action*
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 546
- *action*

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax (loopback | no-loopback);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces interface-range *name* ether-options]

For QFX Series and EX Series:

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],

For SRX Series Devices and vSRX:

[edit interfaces *interface-name* redundant-ether-options]

Release Information Statement introduced before Junos OS Release 7.4 for MX Series.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement modified in Junos OS Release 9.2 for the SRX Series.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the **loopback** statement is configured at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level, an NDP address cannot be configured at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number* family inet6 address]** hierarchy level.

Default By default, loopback is disabled.

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

Related Documentation

- Configuring Ethernet Loopback Capability on page 18
- *Understanding Interfaces*

loopback (Local and Remote)

Syntax loopback (local | remote);

Hierarchy Level [edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers.
Statement introduced in Junos OS Release 15.1F6 and 16.1R2 for PTX3000 routers.

Description Enables local loopback and enables remote loopback. This allows you to test the transceiver cable connection from the far end to the retimer interface without changing the cable.

Options **local**—Enables local loopback

remote—Enables remote loopback

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

Related Documentation

- *Configuring Ethernet Loopback Capability*

loopback-tracking

Syntax loopback-tracking;

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

Release Information Statement introduced in Junos OS Release 14.2.

Description Enables loopback tracking on Ethernet interfaces. When loopback tracking is enabled and the Ethernet Operation, Administration, and Management (OAM) link-fault management process (lfmd) detects its own generated packets on an interface, it marks the interface as down. When the loopback issue resolves, the interface is brought back up.

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 686
- Enabling IEEE 802.3ah OAM Support on page 691

loss-priority

Syntax loss-priority (high | low);

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the packet loss priority value.

Options **high**—Packet has high loss priority.

low—Packet has low loss priority.

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

Related Documentation

- Specifying an Output Priority Map on page 541

mac

Syntax mac *mac-address*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the MAC address of the interface.

Use this statement at the **[edit interfaces ... ps0]** hierarchy level to configure the MAC address for a pseudowire logical device that is used for subscriber interfaces over point-to-point MPLS pseudowires.

Options *mac-address*—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

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

Related Documentation

- Configuring the MAC Address on the Management Ethernet Interface on page 27
- *Configuring a Pseudowire Subscriber Logical Interface Device*

mac (IRB)

Syntax mac *mac-address*;

Hierarchy Level [edit interfaces *irb* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify the MAC address of the IRB interface in devices that have Modular Port Concentrator (MPC) cards.

Options **mac-address**— Specify the MAC address as six hexadecimal bytes in one of the following hexadecimal formats: *nnnn:nnnn:nnnn*or *nn:nn:nn:nn:nn:nn.*

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

Related Documentation

- Example: Configuring the MAC Address of an IRB Interface on page 327

mac-address (Accept Source Mac)

Syntax

mac-address *mac-address* policer;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* accept-source-mac],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* accept-source-mac]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, specify a remote MAC address on which to count incoming and outgoing packets.

Options

mac-address—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**.

policer—MAC policer. For more information, see policer (MAC).

Required Privilege Level

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

Related Documentation

- Configuring MAC Address Filtering on page 544

mac-address (MACsec)

Syntax mac-address *mac-address*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name* id]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specify a MAC address to enable MACsec using static secure association key (SAK) security mode. The **mac-address** variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.

If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the **mac-address**.

If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the **mac-address**.

You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.

Default No MAC address is specified in the secure channel, by default.

Options ***mac-address***—The MAC address, in six groups of two hexadecimal digits.

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

mac-learn-enable

Syntax

mac-learn-enable;

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile]
[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Support for statement under the **[edit interfaces ae*x* aggregated-ether-options ethernet-switch-profile]** hierarchy introduced in Junos OS Release 15.1.

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, for 100-Gigabit Ethernet Type 5 PIC with CFP, and for MPC3E, MPC4E, MPC5E, MPC5EQ, and MPC6E MPCs, configure dynamic learning of the source and destination MAC addresses. By default, the interface is not allowed to dynamically learn source and destination MAC addresses.

To disable dynamic learning of the source and destination MAC addresses after it has been configured, you must delete **mac-learn-enable** from the configuration.

MPCs support MAC address accounting for an individual interface or an aggregated Ethernet interface member link only after the interface has received traffic from the MAC source. If traffic is only exiting an interface, the MAC address is not learned and MAC address accounting does not occur.

Required Privilege Level

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

Related Documentation

- Configuring MAC Address Filtering on page 544
- Configuring MAC Address Accounting on page 531

mac-radius (MX Series in Enhanced LAN Mode)

Syntax mac-radius <flap-on-disconnect> <restrict>;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure MAC RADIUS authentication for specific interfaces. MAC RADIUS authentication allows LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the switch consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN.

If MAC RADIUS is configured, the switch first tries to get a response from the host for 802.1X authentication. If the host is unresponsive, the switch attempts to authenticate using MAC RADIUS.

To restrict authentication to MAC RADIUS only, use the **restrict** option. In restrictive mode, all 802.1X packets are eliminated and the attached device on the interface is considered a nonresponsive host.

Options **flap-on-disconnect**—(Optional) When the RADIUS server sends a disconnect message to a supplicant, the switch resets the interface on which the supplicant is authenticated. If the interface is configured for multiple supplicant mode, the switch resets all the supplicants on the specified interface. This option takes effect only when the **restrict** option is also set.

restrict—(Optional) Restricts authentication to MAC RADIUS only. When **mac-radius restrict** is configured the switch drops all 802.1X packets. This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface, and eliminates the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.

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

mac-validate

Syntax mac-validate (loose | strict);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable IP and MAC address validation for static Ethernet and IP demux interfaces.

Options **loose**—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

strict—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

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

Related Documentation

- MAC Address Validation on Static Ethernet Interfaces Overview on page 239
- *Configuring an IP Demultiplexing Interface*
- *Configuring a VLAN Demultiplexing Interface*

macsec (MX Series)

Syntax

```
macsec {
  connectivity-association connectivity-association-name {
    cipher-suite encryption-algorithm-name;
      exclude-protocol protocol-name;
    pre-shared-key-chain macsec-pre-shared-key-chain-name
    include-sci;
    mka {
      must-secure;
      key-server-priority priority-number;
      transmit-interval interval;
    }
    no-encryption;
    offset (0|30|50);
    pre-shared-key {
      cak hexadecimal-number;
      ckn hexadecimal-number;
    }
    replay-protect{
      replay-window-size number-of-packets;
    }
    secure-channel secure-channel-name {
      direction (inbound | outbound);
      encryption ;
      id {
        mac-address mac-address;
        port-id port-id-number;
      }
      offset (0|30|50);
      security-association security-association-number {
        key key-string;
      }
    }
    security-mode security-mode;
  }
  interfaces interface-name {
    connectivity-association connectivity-association-name;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Configure Media Access Control Security (MACsec) on MX Series routers.

Options 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 Media Access Control Security (MACsec) on MX Series Routers*

major-ring-name

Syntax major-ring-name *name*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify the name of major ring to which the sub-ring node is interconnected.

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

Related Documentation
- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

manual switch

Syntax request protection-group ethernet-aps manual-switch md <md> ma <ma>

Hierarchy Level [edit protocols protection-group ethernet-aps]

Description Forces traffic to switch from the active path to the alternate path, even in the absence of a failure on the working path. If the working path is the active path, traffic will be switched to the protection path. If the protection path is the active path, traffic will be switched to the protection path.

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

Related Documentation
- Ethernet Automatic Protection Switching Overview on page 213

master-only

Syntax master-only;

Hierarchy Level [edit groups rex interfaces (fxp0 | em0) unit *logical-unit-number* family *family* address],
[edit groups rex logical-systems *logical-system-name* interfaces fxp0 unit *logical-unit-number* family *family* address],
[edit interfaces (fxp0 | em0) unit *logical-unit-number* family *family* address],
[edit logical-systems *logical-system-name* interfaces fxp0 unit *logical-unit-number* family *family* address]

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

Description Configure the IP address to be used when the Routing Engine is the current master.

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

Related Documentation

- Configuring a Consistent Management IP Address on page 26
- *CLI User Guide*

max-sessions (PPPoE Service Name Tables)

Syntax max-sessions *number*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the maximum number of active PPPoE sessions using either static or dynamic PPPoE interfaces that the router can establish with the specified named service, **empty** service, or **any** service entry in a PPPoE service name table. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

The router uses the **max-sessions** value for a PPPoE service name table entry in conjunction with the **max-sessions** value configured for the PPPoE underlying interface, and with the maximum number of PPPoE sessions supported on your router. If your configuration exceeds any of these maximum session limits, the router is unable to establish the PPPoE session.

Options *number*—Maximum number of active PPPoE sessions that the router can establish with the specified PPPoE service name table entry, in the range 1 to the platform-specific maximum PPPoE sessions supported for your router. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform.

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

Related Documentation

- *Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name*
- *Configuring PPPoE Service Name Tables*
- *PPPoE Maximum Session Limit Overview*
- *Configuring an Interface Set of Subscribers in a Dynamic Profile*
- *Subscriber Interfaces and PPPoE Overview*

max-sessions-vsa-ignore (Static and Dynamic Subscribers)

Syntax

max-sessions-vsa-ignore;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information

Statement introduced in Junos OS Release 11.4.

Description

Configure the router to ignore (clear) the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143], and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement. The PPPoE maximum session value specifies the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

Default

If you do not include the **max-sessions-vsa-ignore** statement, the maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

Required Privilege Level

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

Related Documentation

- *Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface*
- *PPPoE Maximum Session Limit Overview*
- *Guidelines for Using PPPoE Maximum Session Limit from RADIUS*
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- *Configuring an Interface Set of Subscribers in a Dynamic Profile*
- *Subscriber Interfaces and PPPoE Overview*

maximum-links

Syntax maximum-links *maximum-links-limit*;

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced in Junos OS Release 11.1 for T Series routers.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.3 for MX Series routers.

Description Configure the maximum links limit for aggregated devices. Note that for MX Series routers, to set a range of 32 or 64 the router must be running in Enhanced IP mode, which is only supported for Trio-based MPCs and multiservice DPCs (MS-DPCs). For more information on Enhanced IP mode, *Network Services Mode Overview*.

For MX series routers and PTX series switches, the option for 64 links is only supported for Junos OS release 12.3 and later.

NOTE: **This statement is not supported on the MX80, MX104, and PTX1000 routers.**

Options *maximum-links-limit*—Maximum links limit for aggregated devices.

Range: 16, 32, 64

NOTE: **On T-Series routers, the maximum-links supported is 32 in an aggregated Ethernet link.**

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

Related Documentation

- *Network Services Mode Overview*
- Configuring Junos OS for Supporting Aggregated Devices on page 129
- Configuring an Aggregated Ethernet Interface on page 108
- *network-services*

maximum-requests

Syntax maximum-requests *times*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the maximum number of retransmission times of an EAPOL Request packet to the client before it times out the authentication session.

Options *times*—Specify the maximum number of retransmission times.

Range: 1 through 10 times

Default: 2 times

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- authenticator on page 1087
- dot1x on page 1124
- interface (IEEE 802.1x) on page 1205

maximum-requests (MX Series in Enhanced LAN Mode)

Syntax maximum-requests *number*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X authentication, configure the maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out.

Default Two retransmission attempts

Options *number*—Number of retransmission attempts.

Range: 1 through 10

Default: 2

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

mc-ae

Syntax

```
mc-ae {
  chassis-id chassis-id;
    events {
      iccp-peer-down;
        force-icl-down;
      prefer-status-control-active;
      }
      init-delay-time seconds;
      mc-ae-id mc-ae-id;
      mode (active-active | active-standby);
      redundancy-group group-id;
      revert-time revert-time;
      status-control (active | standby);
      switchover-mode (non-revertive |revertive);
    }
```

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options],
[edit logical-systems *logical-system-name interfaces* ae*X* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.
events statement introduced in Junos OS Release 11.4R4 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series. Only the **chassis-id**, **mc-ae-id**, **mode active-active**, and **status-control (active | standby)** options are supported on QFX Series devices.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
prefer-status-control-active statement introduced in Junos OS Release 13.2R1 for EX Series switches.
init-delay-time *seconds* statement introduced in Junos OS Release 13.2R3 for EX Series switches.
switchover-mode and **revert-time** statements introduced in Junos OS Release 13.3.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Enable multichassis link aggregation groups (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

Options

chassis-id—Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links.

Values: 0 or 1

events—Specify an action if a specific MC-LAG event occurs.

iccp-peer-down—Specify an action if the ICCP peer of this node goes down.

force-icl-down—If the node’s ICCP peer goes down, bring down the interchassis-link logical interface.

prefer-status-control-active—Specify that the node configured as **status-control active** become the active node if the peer of this node goes down.

When ICCP goes down, you can use this keyword to make a mc-lag PE to become the active PE. For example, if you want mc-lag PE1 to be Active on ICCP down, then configure this keyword in PE1. It is not recommended to configure this keyword in both the mc-lag PEs.

NOTE: The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **status-control active** configuration goes down, the router or switch with the **status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lacp periodic *interval*** statement at the **[edit interface *interface-name* aggregated-ether-options]** hierarchy level as **slow** or configure the **detection-time threshold** statement at the **[edit protocols iccp peer liveness-detection]** hierarchy level as less than 3 seconds.

init-delay-time *seconds*—To minimize traffic loss, specify the number of seconds in which to delay bringing the multichassis aggregated Ethernet interface back to the up state when you reboot an MC-LAG peer.

mc-ae-id *mc-ae-id*—Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

Range: 1 through 65,535

mode (active-active | active-standby)—Specify whether the MC-LAG is in active-active or active-standby mode.

NOTE: You can configure IPv4 (inet) and IPv6 (inet6) addresses on **mc-ae** interfaces when the **active-standby** mode is configured.

redundancy-group ***group-id***—Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.

Range: 1 through 4,294,967,294

revert-time—Wait interval (in minutes) before the switchover to the preferred node is performed when the **switchover-mode** is configured as revertive.

Range: 1 through 10

status-control (active | standby)—Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.

switchover-mode (non-revertive | revertive)—Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.

NOTE: For **revertive** mode to automatically switch over to the preferred node, the **status-control** statement should be configured as **active**.

init-delay-time ***seconds***—To minimize traffic loss, specify the number of seconds by which to delay bringing the multichassis aggregated Ethernet (mc-ae) interface back to the up state when you reboot an MC-LAG peer.

Required Privilege Level

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

Related Documentation

- *Active-Active Bridging and VRRP over IRB Functionality Overview*
- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Multichassis Link Aggregation on EX Series Switches*
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- *Example: Configuring Multichassis Link Aggregation in Active-Active Mode*
- *Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces*

minimum-bandwidth (aggregated Ethernet)

Syntax minimum-bandwidth bw-unit *unit* bw-value *value*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options]

Release Information Statement introduced before Junos OS Release 14.1R1 and 14.2 for MX Series.

Description Configure the minimum bandwidth unit for an aggregated Ethernet bundle as bps, Gbps, Kbps, or Mbps and the bandwidth value from 1 through 128,000.

Options ***unit***—Minimum bandwidth unit for the aggregated Ethernet bundle as bps, Gbps, Kbps, or Mbps.

value—Minimum bandwidth value from 1 through 128,000.

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

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 102
- Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 117
- Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles on page 124

minimum-links

Syntax (SRX, MX, T, M, EX, QFX Series, EX4600, Qfabric System)

minimum-links *number*;

Hierarchy Level (EX Series)

[edit interfaces ae*x* aggregated-ether-options],
[edit interfaces ae*x* aggregated-sonet-options],
[edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *range* aggregated-ether-options],
[edit interfaces interface-range *range* aggregated-sonet-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Hierarchy Level (QFX Series)

[edit interfaces ae*x* aggregated-ether-options]

Release Information

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

Description

For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

Options

number—Number of links.

Range: On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On EX8200 switches, the range of valid values for minimum-links number is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On EX4600, QFX Series and Q Fabric Systems , the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

Default: 1

Required Privilege Level

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

Related Documentation

- Configuring Aggregated Ethernet Minimum Links on page 137
- *Configuring Aggregated SONET/SDH Interfaces*
- *Configuring Aggregated Ethernet Links (CLI Procedure)*

- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Configuring Link Aggregation*

mixed-rate-mode

Syntax mixed-rate-mode;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number* mixed-rate-mode],
[edit chassis *lcc number* fpc *slot-number* pic *pic-number* mixed-rate-mode] (Routing Matrix)

Release Information Statement introduced in Junos OS Release 13.3.

Description Configure the mixed-rate mode for the 24-port 10 Gigabit Ethernet PIC (PF-24XGE-SFPP) only.

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

Related Documentation

- Modes of Operation of 10-Gigabit Ethernet PICs on page 403
- Configuring Mixed-Rate Mode Operation on page 408

mka (MX Series)

Syntax

```
mka {
  must-secure;
  key-server-priority priority-number;
  transmit-interval interval;
}
```

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specify parameters for the MACsec Key Agreement (MKA) protocol.

Options 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 Media Access Control Security (MACsec) on MX Series Routers*

must-secure (MX Series)

Syntax must-secure;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* mka]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies that all traffic travelling on the MACsec-secured link must be MACsec-secured to be forwarded onward.

When the **must-secure** option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.

When the **must-secure** option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do no support MACsec is forwarded through the network.

The **must-secure** option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn’t support MACsec can continue to send and receive traffic over the network—provided the **must-secure** option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.

Default The **must-secure** option is disabled.

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

mtu

Syntax

mtu *bytes*;

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit interfaces interface-range *name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* protocols l2circuit local-switching interface *interface-name* backup-neighbor *address*],
[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name* backup-neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls],
[edit protocols l2circuit local-switching interface *interface-name* backup-neighbor *address*],
[edit protocols l2circuit neighbor *address* interface *interface-name*]
[edit protocols l2circuit neighbor *address* interface *interface-name* backup-neighbor *address*],
[edit routing-instances *routing-instance-name* protocols l2vpn interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols vpls],
[edit logical-systems *name* protocols ospf area *name* interface],
[edit logical-systems *name* routing-instances *name* protocols ospf area *name* interface],
[edit protocols ospf area *name* interface],
[edit routing-instances *name* protocols ospf area *name* interface]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Support at the **[set interfaces interface-name unit logical-unit-number family ccc]** hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.
Statement introduced in Junos OS 17.3R1 Release for MX Series Routers.

Description

Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named irb or vlan, respectively).

CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.

NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the **flexible-vlan-tagging** statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
- In case of Layer 2 IFL configured with the **vlan-tagging** statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.

NOTE:

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the **mtu** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** or **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level.
  - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
  - If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.
- Starting in Release 14.2, MTU for IRB interfaces is calculated by removing the Ethernet header overhead (**6(DMAC)+6(SMAC)+2(EtherType)**), and the MTU is a minimum of the two values:
  - Configured MTU
  - Associated bridge domain's physical or logical interface MTU
    - For Layer 2 logical interfaces configured with **flexible-vlan-tagging**, IRB MTU is calculated by including 8 bytes overhead (**SVLAN+CVLAN**).
    - For Layer 2 logical interfaces configured with **vlan-tagging**, IRB MTU is calculated by including single VLAN 4 bytes overhead.

  NOTE: Changing the Layer 2 logical interface option from **vlan-tagging** to **flexible-vlan-tagging** or vice versa adjusts the logical interface MTU by 4 bytes with the existing MTU size. As a result, the Layer 2 logical interface is deleted and re-added, and the IRB MTU is re-computed appropriately.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configuring the Media MTU*.

Options *bytes*—MTU size.

Range: 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers), 256 through 9500 bytes (Junos OS 16.1R1 for MX Series routers)

NOTE: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:

- MPC1
- MPC2
- MPC2E
- MPC3E
- MPC4E
- MPC5E
- MPC6E

Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

Required Privilege Level

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

Related Documentation

- *Configuring the Media MTU*
- *Configuring the MTU for Layer 2 Interfaces*
- *Setting the Protocol MTU*

multicast-router-interface (IGMP Snooping)

Syntax multicast-router-interface;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit protocols igmp-snooping vlan (all | *vlan-name*) interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]
[edit protocols igmp-snooping vlan *vlan-name* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.1 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Statically configure the interface as an IGMP snooping multicast-router interface—that is, an interface that faces toward a multicast router or other IGMP querier.

NOTE: If the specified interface is a trunk port, the interface becomes a multicast-routing device interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast routing device interface, even if the interface is configured as a multicast routing device interface only for IGMP snooping.

Configure an interface as a bridge interface toward other multicast routing devices.

Default Disabled. If this statement is disabled, the interface drops IGMP messages it receives.

The interface can either be a host-side or multicast-routing device interface.

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

Related Documentation

- *Example: Configuring IGMP Snooping*
- *IGMP Snooping in MC-LAG Active-Active Mode*
- *host-only-interface*

multi-chassis-protection

Syntax

```
multi-chassis-protection {
  peer a.b.c.d {
    interface interface-name;
    }
  }
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, you can use this statement under the physical interface level to reduce the configuration at the logical interface level if the following assumption exists:

If there are n +1 logical interfaces under **ae0**, from **ae0.0** through **ae0.n**, there will be n + 1 logical interfaces under **ge-0/0/0** as well, from **ge-0/0/0.0** through **ge-0/0/0.n**, and each **ge-0/0/0** logical interface will be a protection link for the **ae0** logical interface.

NOTE: **A bridge domain cannot have MC-AE logical interfaces which belong to different redundancy groups.**

If the Inter-Chassis Control Protocol (ICCP) connection is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer.

The remaining statements are explained separately. See CLI Explorer.

Options **interface** ***interface-name***—Specify the interface: **interface** ***interface-name-fpc/pic/port***

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

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- Configuring Aggregated Ethernet Link Protection on page 148
- Example: Configuring Aggregated Ethernet Link Protection on page 150
- peer on page 1301

negotiate-address

Syntax negotiate-address;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.

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

Related Documentation

- *Configuring IPCP Options for Interfaces with PPP Encapsulation*
- address on page 1071
- unnumbered-address (PPP) on page 1429
- *Junos OS Administration Library*

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 remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 686

no-adaptive

Syntax no-adaptive;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options load-balance]

Release Information Statement introduced in Junos OS Release 13.2R3.

Description Configure no-adaptive on the aggregated Ethernet bundle to remove the adaptive that is configured to address the traffic imbalance.

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

Related Documentation
- Understanding Aggregated Ethernet Load Balancing on page 158

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 695

no-encryption (MACsec for MX Series)

Syntax no-encryption;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Disables MACsec encryption for a connectivity association that is configured to enable MACsec using static connectivity association key (CAK) or dynamic security mode.

You can enable MACsec without enabling encryption. If a connectivity association that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the packet, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic does not represent a security threat.

This command is used to disable encryption when MACsec is configured using static CAK or dynamic security mode only. When MACsec is configuring using static secure association key (SAK) security mode, the encryption setting is managed in the secure channel using the **encryption** configuration statement.

Default MACsec encryption is enabled if MACsec is enabled using static CAK or dynamic security mode.

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

no-auto-mdix

Syntax no-auto-mdix;

Hierarchy Level [edit interface ge-*fpc/port/pic* gigether-options]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description Disable the Auto MDI/MDIX feature.

MX Series routers with Gigabit Ethernet interfaces automatically detect MDI and MDIX port connections. Use this statement to override the default setting. Remove this statement to return to the default setting.

Default Auto MDI/MDIX is enabled by default.

Options There are no options for this statement.

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

Related Documentation

- Ethernet Interfaces Overview on page 3
- **gigether-options on page 1180**.

no-gratuitous-arp-request

Syntax no-gratuitous-arp-request;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Ethernet interfaces and pseudowire logical interfaces, do not respond to gratuitous ARP requests.

Default Gratuitous ARP responses are enabled on all Ethernet interfaces.

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

Related Documentation

- Configuring Gratuitous ARP on page 19

no-keepalives

Syntax

no-keepalives;

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Disable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation. The default keepalive interval is 10 seconds.

For ATM2 IQ interfaces only, you can disable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

Required Privilege Level

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

Related Documentation

- *Configuring Keepalives*
- Disabling the Sending of PPPoE Keepalive Messages on page 359
- *Configuring Frame Relay Keepalives*

no-mac-table-binding (802.1X for MX Series in Enhanced LAN Mode)

Syntax no-mac-table-binding;

Hierarchy Level [edit protocols authentication-access-control]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X authentication, disable the removal of the session from the authentication session table when the MAC address ages out of the Ethernet switching table.

Default Not enabled

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

Related Documentation

no-native-vlan-insert

Syntax

no-native-vlan-insert;

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 17.1R1.

Description

Send traffic without the native VLAN ID (*native-vlan-id*) to the remote end of the network if untagged traffic is received.

If this statement is not configured, then *native-vlan-id* is added to untagged traffic. But if this statement is configured, then *native-vlan-id* is not added to untagged traffic.

NOTE:

- This feature works only on MX Series routers with MPCs/MICs. Configuring this statement on MX Series routers with DPCs results is no behavioral change. However, if you configure the statement on aggregated Ethernet (ae) interfaces with logical interfaces across MPCs/MICs and DPCs, then the MPCs/MICs and DPCs behave differently.
- In the egress direction, this feature is disrupted by VLAN normalization. Because of normalization, the egress interface cannot distinguish between untagged traffic and tagged traffic. And untagged traffic is sent out with *native-vlan-id*. Consider this while configuring both VLAN normalization and new *native-vlan-id* statement.

  There will be a problem with ingress firewall filter if filter term includes *native-vlan-id*. With **no-native-vlan-insert** statement configured, *native-vlan-id* will not be inserted to untagged traffic. So, firewall filter term will not match with untagged traffic. But if incoming traffic have VLAN ID which is equal to *native-vlan-id*, then firewall filter term will match and firewall will work.
- When this feature is used with AE, all sub-interfaces of AE should be in same type of FPC.

Default

By default, *native-vlan-id* is inserted to untagged traffic. That is, if this statement is not configured, then *native-vlan-id* is inserted to untagged traffic.

Required Privilege Level

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

Related Documentation

- *Configuring Mixed Tagging Support for Untagged Packets*
- Configuring Access Mode on a Logical Interface on page 261

- *Configuring the Native VLAN Identifier on Switches With ELS Support (CLI Procedure)*
- *Understanding Bridging and VLANs on Switches*
- **flexible-vlan-tagging on page 1169**
- *native-vlan-id*
- *Understanding Q-in-Q Tunneling and VLAN Translation*
- Sending Untagged Traffic Without VLAN ID to Remote End on page 249

no-pre-classifier

Syntax no-pre-classifier;

Hierarchy Level [edit chassis fpc *n* pic *n*]

Release Information Statement introduced in Junos OS Release 10.4.

Description Specify disabling the control queue for all ports on the 10-Gigabit Ethernet LAN/WAN PIC. Deleting this configuration re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

NOTE: For the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), the control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. If the transit control traffic crosses this limit, then it can cause drops on locally terminating control traffic, causing flap of protocols such as BGP and OSPF. To avoid the control traffic being dropped, configure the **no-pre-classifier** statement to disable the control queue.

Default The **no-pre-classifier** statement is not configured and the control queue is operational.

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

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 395
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 404

no-reauthentication (MX Series in Enhanced LAN Mode)

Syntax no-reauthentication;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X authentication, disables reauthentication.

Default Not disabled

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

no-send-pads-ac-info

Syntax no-send-pads-ac-info;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 12.2.

Description Prevent the router from sending the AC-Name and AC-Cookie tags in the PPPoE Active Discovery Session (PADS) packet. When you configure this statement, it affects PADS packets sent on all PPPoE interfaces configured on the router after the command is issued; it has no effect on previously created PPPoE interfaces. By default, the AC-Name and AC-Cookie tags are transmitted in the PADS packet, along with the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags.

NOTE: In Junos OS Release 12.1 and earlier, only the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags are contained in the PADS packet by default. The AC-Name and AC-Cookie tags are not transmitted in the PADS packet by default.

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 PPPoE Access Concentrator Tags in PADS Packets*

no-send-pads-error

Syntax no-send-pads-error;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 12.3.

Description Discard PADR messages to prevent transmission of PADS control packets with AC-System-Error tags.

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

Related Documentation

- *Discarding PADR Messages to Accommodate Abnormal CPE Behavior*

non-revertive (Interfaces)

Syntax non-revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

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

Related Documentation

- **link-protection on page 1227**
- Configuring Aggregated Ethernet Link Protection on page 148
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

non-revertive

Syntax non-revertive;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Enable nonrevertive operation where traffic is allowed to use the RPL if it has not failed, even after a switch condition has cleared. The default mode of operation is revertive.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

non-vc-mode

Syntax non-vc-mode;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Configure a node on the sub-ring to operate in non-virtual channel mode. If this option is enabled then all the nodes in the sub-ring are configured with this option. Also, the **non-vc-mode** option should be used with care and only for open rings. Using this option for closed rings creates loops for RAPS control messages.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

node-id

Syntax node-id *mac-address*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.

Description For EX Series switches and QFX Series switches, node-id is not configurable.

For MX Series routers, optionally specify the MAC address of a node in the protection group. If this statement is not included, the router assigns the node's MAC address.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

offset (MX Series)

Syntax

offset (0 |30 | 50);

Hierarchy Level

[edit security macsec connectivity-association *connectivity-association-name*]
[edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name*]

Release Information

Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description

Specifies the number of octets in an Ethernet frame that are sent in unencrypted plain-text when encryption is enabled for MACsec.

Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

You configure the **offset** in the **[edit security macsec connectivity-association *connectivity-association-name*]** hierarchy when you are enabling MACsec using static connectivity association key (CAK) or dynamic security mode.

You configure the **offset** in the **[edit security macsec connectivity-association *connectivity-association-name* secure-channel*secure-channel-name*]** hierarchy when you are enabling MACsec using static secure association key (SAK) security mode.

Default

0

Options

0—Specifies that no octets are unencrypted. When you set the offset to 0, all traffic on the interface where the connectivity association or secure channel is applied is encrypted.

30—Specifies that the first 30 octets of each Ethernet frame are unencrypted.

NOTE: In IPv4 traffic, setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 30, therefore, is typically used when a feature needs this information to perform a task on IPv4 traffic.

50—Specified that the first 50 octets of each Ethernet frame are unencrypted.

NOTE: In IPv6 traffic, setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 50, therefore, is typically used when a feature needs this information to perform a task on IPv6 traffic.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

oam

Syntax

```
oam {
  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 (100ms | 10m | 10ms | 10s | 1m | 1s);
            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;
                  }
                }
              }
            }
          }
        }
        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;
            }
          }
        }
      }
    }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers and PTX Series Packet Transport Routers, provide IEEE 802.3ah Operation, Administration, and Maitenance (OAM) support.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 686

optics-options

Syntax

```
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
    }
  tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number);
  tx-power dbm;
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
```

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
alarm option and **warning** options introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement and **tx-power** option introduced in Junos OS Release 13.2 for PTX Series routers.
tca option introduced in Junos OS Release 14.2 for PTX Series routers.

Description

For 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475

otn-options

Syntax

```
otn-options {
  bytes (otn-options) transmit-payload-type value;
  fec (efec | gfec | gfec-sdfec | none );
  (is-ma | no-is-ma);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  (local-loopback | no-local-loopback);
  (odu-ttim-action-enable | no-odu-ttim-action-enable);
  (otu-ttim-action-enable | no-otu-ttim-action-enable);
  odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    number-of-frames value;
    (no-start-measurement | start-measurement;
  }
  odu-signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  (prbs | no-prbs);
  preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
    odu-backward-frr-enable | no-odu-backward-frr-enable;
    odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable;
  }
  rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    oc192;
    otu4;
    (pass-through | no-pass-through);
  }
  signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  tca tca-identifier (enable-tca | no-enable-tca) (threshold number | threshold-24hrs number);
  transport-monitoring;
  trigger trigger-identifier;
  tti tti-identifier;
}
```

Hierarchy Level

[edit interfaces ge-*fpc/pic/port*]
[edit interfaces xe-fpc/pic/port]
[edit interfaces et-fpc/pic/port]

Release Information

Statement introduced in Junos OS Release 9.4.

bytes, **is-ma**, **local-loopback**, **no-is-ma**, **no-local-loopback**, **no-odu-ttim-action-enable**, **no-otu-ttim-action-enable**, **no-prbs**, **odu-delay-management**, **odu-ttim-action-enable**,

otu-ttim-action-enable, **prbs**, **preemptive-fast-reroute**, and **signal-degrade** statements introduced in Junos OS Release 13.2 for PTX Series routers.
oc192 statement introduced in Junos OS Release 13.3R3 for MX Series routers.
odu-signal-degrade, **odu-backward-frr-enable | no-odu-backward-frr-enable**, **odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable** statements introduced in Junos OS Release 14.1R2 and 14.2 for P2-100GE-OTN PIC in PTX5000 routers.
tca option introduced in Junos OS Release 14.2 for PTX Series routers.

Description Specify the Ethernet optical transport network (OTN) interface and options.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 465
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- Configuring 100-Gigabit DWDM OTN PICs on page 482

output-policer

Syntax

output-policer *policer-name*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* layer2-policer],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Apply a single-rate two-color policer to the Layer 2 output traffic at the logical interface. The **output-policer** and **output-three-color** statements are mutually exclusive.

Options

policer-name—Name of the single-rate two-color policer that you define at the **[edit firewall]** hierarchy level.

Required Privilege Level

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- Example: Configuring Gigabit Ethernet Policers on page 544
- **input-policer on page 1199**
- **input-three-color on page 1201**
- **layer2-policer on page 1219**
- **logical-interface-policer on page 1238**
- **output-three-color on page 1293**

output-priority-map

Syntax

```
output-priority-map {
  classifier {
    premium {
      forwarding-class class-name {
        loss-priority (high | low);
      }
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the output policer priority map to be applied to outgoing frames on this interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Specifying an Output Priority Map on page 541
- input-priority-map on page 1200

output-three-color

Syntax

output-three-color *policer-name*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* layer2-policer]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* layer2-policer]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Apply a single-rate or two-rate three-color policer to the Layer 2 output traffic at the logical interface. The **output-three-color** and **output-policer** statements are mutually exclusive.

Options

policer-name—Name of the single-rate or two-rate three-color policer.

Required Privilege Level

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

Related Documentation

- *Two-Color and Three-Color Policers at Layer 2*
- *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*
- Example: Configuring Gigabit Ethernet Policers on page 544
- **input-three-color on page 1201**
- **input-policer on page 1199**
- **layer2-policer on page 1219**
- **logical-interface-policer on page 1238**
- **output-policer on page 1291**

output-vlan-map (Aggregated Ethernet)

Syntax

```
output-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 8.2.
Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

Description

Define the rewrite profile to be applied to outgoing frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet inferfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP..

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 560
- **input-vlan-map (Aggregated Ethernet) on page 1202**

output-vlan-map

Syntax

```
output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements added in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For EX Series switches, defines the rewrite operation to be applied to outgoing frames.

For MX Series routers and NFX Series devices' Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, defines the rewrite operation to be applied to outgoing frames on this logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 560
- input-vlan-map on page 1203

pado-advertise

Syntax pado-advertise;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable named services configured in PPPoE service name tables to be advertised in PPPoE Active Discovery Offer (PADO) control packets. By default, advertisement of named services in PADO packets is disabled.

NOTE: If you enable advertisement of named services in PADO packets, make sure the number and length of of all advertised service entries does not exceed the maximum transmission unit (MTU) size of the PPPoE underlying interface.

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Enabling Advertisement of Named Services in PADO Control Packets*

passive-monitor-mode

Syntax passive-monitor-mode;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Monitor packet flows from another router. If you include this statement in the configuration, the interface does not send keepalives or alarms, and does not participate actively on the network.

This statement is supported on ATM, Ethernet, and SONET/SDH interfaces. For more information, see *ATM Interfaces Feature Guide for Routing Devices*.

For ATM and Ethernet interfaces, you can include this statement on the physical interface only.

For SONET/SDH interfaces, you can include this statement on the logical interface only.

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

Related Documentation

- *Enabling Passive Monitoring on ATM Interfaces*
- Passive Monitoring on Ethernet Interfaces Overview on page 29
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*
- *multiservice-options*
- *Junos OS Services Interfaces Library for Routing Devices*

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 for MX, M, T, ACX, Series routers, SRX Series firewalls, and EX Series Switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches.

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: For MX, M, T, ACX, Series routers, SRX Series firewalls and EX Series switches – 100 through 1000 milliseconds

Default: For EX Series switches –1000 milliseconds

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

Related Documentation

- Configuring the OAM PDU Interval on page 693
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

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 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

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 694
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

per-flow (Aggregated Ethernet Interfaces)

Syntax per-flow;

Hierarchy Level [edit interfaces ae*X unit logical-unit-number* forwarding-options load-balance-stateful]

Release Information Statement introduced in Junos OS Release 13.2R1.

Description Enable the mechanism to perform an even, effective distribution of traffic flows across member links of an aggregated Ethernet interface (**ae**) bundle on MX Series routers with MPCs, except MPC3Es and MPC4Es. When multiple flows are transmitted out of an **ae** interface, the flows must be distributed across the different member links evenly to enable an effective and optimal load-balancing behavior. To obtain a streamlined and robust method of load-balancing, the member link of the aggregated Ethernet interface bundle that is selected each time for load balancing plays a significant part.

Options **per-flow**—Enable the stateful load-distribution mechanism per traffic flow on an aggregated Ethernet interface.

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

Related Documentation

- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces on page 181

peer

Syntax

peer *a.b.c.d* {
  interface *interface-name*;
}

Hierarchy Level

[edit interfaces *interface-name* multi-chassis-protection]

Release Information

Statement introduced in Junos OS Release 11.1.

Description

For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, use the **multi-chassis-protection** statement under the physical interface level to reduce the configuration at the logical interface level. If the interchassis control protocol connection (ICCP) is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement. You must also specify the peer's physical interface.

Options

a.b.c.d—Specify the IP address of the peer.

interface *interface-name*—Specify the peer's physical interface: interface *interface-name-fpc/pic/port*

Required Privilege Level

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

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- Configuring Aggregated Ethernet Link Protection on page 148
- Example: Configuring Aggregated Ethernet Link Protection on page 150
- **multi-chassis-protection on page 1269**

periodic

List of Syntax

Syntax (EX Series) on page 1302
Syntax (QFX Series) on page 1302

Syntax (EX Series)

periodic *interval*;

Syntax (QFX Series)

periodic (fast | slow);

Hierarchy Level (EX Series)

[edit interfaces ae*x* aggregated-ether-options lacp],
[edit interfaces interface-range *name* aggregated-ether-options lacp]

Hierarchy Level (QFX Series)

[edit interfaces ae*x* aggregated-ether-options lacp]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description

For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

Options

interval—Interval for periodic transmission of LACP packets.

- **fast**—Transmit packets every second.
- **slow**—Transmit packets every 30 seconds.

Default: fast

Required Privilege Level

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

Related Documentation

- Configuring LACP for Aggregated Ethernet Interfaces on page 140
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- *Junos OS Network Interfaces Library for Routing Devices*

policer (CFM Firewall)

Syntax

```
policer cfm-policer {
  if-exceeding {
    bandwidth-limit 8k;
    burst-size-limit 2k;
  }
  then discard;
}
```

Hierarchy Level [edit firewall]

Release Information Statement introduced in Junos OS Release 10.0.

Description Attach an explicit policer to CFM sessions.

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

Related Documentation

- Configuring Rate Limiting of Ethernet OAM Messages on page 655
- **policer (CFM Global) on page 1039**
- **policer (CFM Session) on page 1040**

policer (CoS)

Syntax

```
policer cos-policer-name {
  aggregate {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
  premium {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ , Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, define a CoS policer template to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits. The premium policer is not supported on MX Series routers or for Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Options

cos-policer-name—Name of one policer to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539

policer (MAC)

Syntax

```
policer {
  input cos-policer-name;
  output cos-policer-name;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* accept-source-mac mac-address *mac-address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* accept-source-mac mac-address *mac-address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure MAC policing.

NOTE:
On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

Options

input *cos-policer-name*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

output *cos-policer-name*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

Required Privilege Level

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

Related Documentation

- Configuring MAC Address Filtering on page 544

pop

Syntax

pop;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

NOTE: On EX4300 switches, pop is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map] hierarchy level.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level

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

Related Documentation

- Removing a VLAN Tag on page 570
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

pop-pop

Syntax

pop-pop;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, and for 10-Gigabit Ethernet SFP interfaces on EX Series switches, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.

Required Privilege Level

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

Related Documentation

- Removing the Outer and Inner VLAN Tags on page 571

pop-swap

Syntax

pop-swap;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2, IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

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

Related Documentation

- Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag on page 572

port-description-type

Syntax

```
port-description-type {
  interface-alias;
  interface-description;
}
```

Hierarchy Level

[edit protocols lldp]
[edit routing- instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 13.3R5, 14.2R2, 14.1R4, and 12.3R9.

Description

For Link Layer Discovery Protocol, configure the value to be used for port description TLV.

Options

interface-alias—Use the *ifAlias* MIB object value to generate the port description TLV.

interface-description—Use the *ifDescr* MIB object value to generate the port description TLV.

Default: By default, **interface-alias** is used for generation of port description TLV. The **interface-alias** value is same as description of an interface configured by statement **set interface *name* description *description***.

NOTE: The LLDP MIB *lldpLocPortDesc* value gets changed depending on the setting of **port-description-type** statement. That is:

- If you configure the **port-description-type interface-alias** statement, the MIB variable *lldpLocPortDesc* displays the value same as that of MIB variable *ifAlias*, which is same as the description of the interface. By default, the MIB variable *lldpLocPortDesc* displays the value same as that of MIB variable *ifAlias*.
- If you configure the **port-description-type interface-description** statement , the MIB variable *lldpLocPortDesc* displays the value of MIB variable *ifDescr*, which is same as that of interface name.

Required Privilege Level

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

Related Documentation

- lldp on page 1232
- show lldp on page 2012
- show lldp neighbors on page 2018

port-id (MACsec for MX Series)

Syntax port-id *port-id-number*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name* id]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specify a port ID in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.

Once the port numbers match, MACsec is enabled for all traffic on the connection.

You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.

Default No port ID is specified.

Options ***port-id-number***—The port ID number.

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

Related Documentation
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

port-priority

Syntax port-priority *priority*;

Hierarchy Level [edit interfaces *interface-name* gigether-options 802.3ad lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define LACP port priority at the interface level.

Options *priority*—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected.

Range: 0 through 65535

Default: 127

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

Related Documentation

- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*

port-id-subtype

Syntax

```
port-id-subtype {
  interface-name;
  locally-assigned;
}
```

Hierarchy Level

[edit protocols lldp]
[edit routing- instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 12.3R1

Description

For Link Layer Discovery Protocol, configure the port ID type, length, and value (TLV).

Options

interface-name—Use the interface name to generate the port ID TLV.

Default: Use the SNMP index of the interface to generate the port ID TLV. This is the default option used to generate port ID TLV.

NOTE: For QFX5200 switches, the default value used to generate port ID TLV in LLDP messages is interface name, not SNMP index.

NOTE: The **show lldp neighbors** command displays the content of the **port ID** TLV received from the peer in the **port Info** field. Changing the configuration of **port-id-subtype** affects the display of the **show lldp neighbors** command on the peer device running Junos OS.

When the value of **port-id-subtype** is set to **locally-assigned**, which is the default value, the **show lldp neighbors** command on the peer device running Junos OS displays the SNMP index as the port information for the local device.

When the value of **port-id-subtype** is set to **interface-name**, the **show lldp neighbors** command on the peer device running Junos OS displays the interface name as the port information for the local device.

The value of the MIB variable instance **lldpLocPortId** depends on the entity that is used to generate the port ID TLV. If the port ID TLV generation is configured to use the interface-name in the command **set port-id-subtype interface-name**, then the value of the MIB variable **lldpLocPortId** instance is the interface name and not the SNMP index.

Required Privilege Level

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

Related Documentation

- **lldp on page 1232**
- Configuring LLDP on page 338
- **show lldp neighbors on page 2018**

pp0 (Dynamic PPPoE)

Syntax

```
pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      aaa-options aaa-options-name;
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
      }
      ignore-magic-number-mismatch;
      initiate-ncp (ip | ipv6 | dual-stack-passive)
      ipcp-suggest-dns-option;
      mru size;
      mtu (size | use-lower-layer);
      on-demand-ip-address;
      pap;
      peer-ip-address-optional;
    }
    family inet {
      unnumbered-address interface-name;
      address address;
      service {
        input {
          service-set service-set-name {
            service-filter filter-name;
          }
          post-service-filter filter-name;
        }
        output {
          service-set service-set-name {
            service-filter filter-name;
          }
        }
      }
      filter {
        input filter-name {
          precedence precedence;
        }
        output filter-name {
          precedence precedence;
        }
      }
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces]

Release Information Statement introduced in Junos OS Release 10.1.

Description Configure the dynamic PPPoE logical interface in a dynamic profile. When the router creates a dynamic PPPoE logical interface on an underlying Ethernet interface configured with PPPoE (**ppp-over-ether**) encapsulation, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical interface.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring a PPPoE Dynamic Profile*
- *Configuring Dynamic Authentication for PPP Subscribers*
- For information about creating static PPPoE interfaces, see Configuring PPPoE on page 353

ppm (Ethernet Switching)

Syntax

```
ppm {
  centralized;
}
```

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced in Junos OS Release 9.4 for MX Series routers.
Statement introduced in Junos OS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.1 for T Series devices.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.

This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** configuration statement in the **[edit routing-options ppm]** statement hierarchy.

Default Distributed PPM processing is enabled for all packets that use PPM.

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

Related Documentation

- *Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)*
- Configuring Distributed Periodic Packet Management on page 197

pppoe-options

Syntax

```
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
  ppp-max-payload ppp-max-payload
}
```

Hierarchy Level

[edit interfaces pp0 unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]
[set interface *ppp interface*unit *logical-unit-number* ppp-max-payload *ppp-max-payload*],

Release Information

Statement introduced before Junos OS Release 7.4.
client Statement introduced in Junos OS Release 8.5.
server Statement introduced in Junos OS Release 8.5.
client Statement introduced in Junos OS Release 15.1X49-D100.

Description

Configure PPP over Ethernet-specific interface properties.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

The PPP-Max-Payload option allows you to override the default behavior of the PPPoE client by providing a maximum size that the PPP payload can support in both sending and receiving directions. The PPPoE server might allow the negotiation of an MRU larger than 1492 octets and the ability to use an MTU larger than 1500 octets.

Required Privilege Level

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

Related Documentation

- *Configuring a PPPoE Interface*

pppoe-underlying-options (Static and Dynamic Subscribers)

Syntax

```
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  direct-connect
  duplicate-protection;
  max-sessions number;
  max-sessions-vsa-ignore;
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds> <lockout-time-max
    maximum-seconds> <filter [aci]>;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (**ppp-over-ether**) encapsulation.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

Required Privilege Level

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

Related Documentation

- Configuring PPPoE on page 353 (for static interfaces)
- *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*
- *Assigning a Service Name Table to a PPPoE Underlying Interface*

preferred-source-address

Syntax preferred-source-address *address*;

Hierarchy Level [edit dynamic-profiles interfaces *interface-name* unit *logical-unit-number* family *family* unnumbered-address *interface-name*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*],

Release Information Statement introduced in Junos OS Release 9.2.
Support for the **$junos-preferred-source-address** and **$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.

Description For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. To configure the preferred source address dynamically, instead of using this statement, you must include the **$junos-preferred-source-address** predefined variable for IPv4 (family inet) addresses or the **$junos-preferred-source-ipv6-address** predefined variable for IPv6 (family inet6) addresses.

Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for IPv4 and IPv6 address families.

NOTE: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the **$junos-routing-instance** predefined variable, you must not configure a preferred source address, whether with the **$junos-preferred-source-address** predefined variable, the **$junos-preferred-source-ipv6-address** predefined variable, or the **preferred-source-address** statement. Configuring the preferred source address in this circumstance causes a commit failure.

Options *address*—Secondary IP address of the donor loopback interface. Alternatively, use the **$junos-preferred-source-address** or the **$junos-preferred-source-ipv6-address** predefined variable to dynamically apply a preferred source address to the unnumbered Ethernet interface.

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

Related Documentation
- *Configuring an Unnumbered Interface*
- *Junos OS Network Interfaces Library for Routing Devices*

- *Junos OS Administration Library*

pre-shared-key (MX Series)

Syntax

```
pre-shared-key {
  cak hexadecimal-number;
  ckn hexadecimal-number;
}
```

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode.

A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the pre-shared keys are successfully verified and exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link.

Default No pre-shared keys exist, by default.

Options 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 Media Access Control Security (MACsec) on MX Series Routers*

premium (Output Priority Map)

Syntax

```
premium {
  forwarding-class class-name {
    loss-priority (high | low);
  }
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Specifying an Output Priority Map on page 541
- input-priority-map on page 1200

premium (Policer)

Syntax

```
premium {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- aggregate (Gigabit Ethernet CoS Policer) on page 1077
- ieee802.1p on page 1188

propagate-tc

Syntax propagate-tc;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Enable topology change propagation from a sub-ring to an interconnected major-ring. By default, topology change propagation is disabled.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

protection-group

Syntax

```
protection-group {
  ethernet-ring ring-name {
    data-channel {
      vlan number
    }
    east-interface {
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    non-revertive;
    wait-to-block-interval number;
    major-ring-name name;
    propagate-tc;
    compatibility-version (1|2);
    ring-id number;
    non-vc-mode;
    dot1p-priority number;
    west-interface {
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
      virtual-control-channel {
        west-interface name;
        east-interface name;
      }
    }
  }
 control-vlan (vlan-id | vlan-name);
    east-interface {
      node-id mac-address;
      control-channel channel-name {
        vlan number;
        interface name interface-name
      }
      interface-none
      ring-protection-link-end;
    }
  }
    control-channel channel-name {
      vlan number;
      interface name interface-name
    }
  }
  data-channel {
    vlan number
```

```
      }
      guard-interval number;
      node-id mac-address;
      restore-interval number;
      ring-protection-link-owner;
        west-interface {
          node-id mac-address;
          control-channel channel-name {
            vlan number;
            interface name interface-name
          }
          interface-none
          ring-protection-link-end;
        }
        control-channel channel-name {
          vlan number;
          interface name interface-name
        }
      }
    }
    guard-interval number;
    restore-interval number;
    traceoptions {
      file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
      flag flag;
    }
  }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure Ethernet ring protection switching.

The statements are explained separately. All statements apply to MX Series routers. EX Series switches do not assign **node-id** and use **control-vlan** instead of **control-channel**.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857
- Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers on page 864
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

protocols

Syntax protocols [inet iso mpls];

Hierarchy Level [edit interfaces *interface-name* unit logical-unit-number family tcc]

Release Information Statement introduced in Junos OS Release 8.3.

Description For Layer 2.5 VPNs on T Series, MX Series, M120, and M320 routers support, configure IS-IS (ISO traffic) or MPLS traffic to traverse a TCC interface. By default, IPv4 (inet) traffic runs on T Series, MX, Series, M120, and M320 routers and over TCC interfaces. You must configure the same traffic type on both ends of the Layer 2.5 VPN.

NOTE: Some platform and FPC combinations can not pass TCC encapsulated ISO traffic. See *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic* for details.

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

Related Documentation

- *Configuring IS-IS or MPLS Traffic for TCC Interfaces*
- *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic*

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 698

ptopo-configuration-maximum-hold-time

Syntax ptopo-configuration-maximum-hold-time *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a time to maintain dynamic topology entries.

Options *seconds*—Time to maintain interval dynamic topology entries.

Default: 300

Range: 1 through 2147483647

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

Related Documentation

- Configuring LLDP on page 338

ptopo-configuration-trap-interval

Syntax ptopo-configuration-trap-interval *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description Configure a time for the period of SNMP trap notifications to the Master Agent to wait regarding changes in topology global statistics.

Options *seconds*—Time for the period of SNMP trap notifications about global statistics. This feature is disabled by default.

Range: 0 through 3600

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

Related Documentation

- Configuring LLDP on page 338

push

Syntax

push;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

NOTE: On EX4300 switches, **push** is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map] hierarchy level.

Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.

You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

If you include the **push** statement in the configuration, you must also include the pop statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level.

Required Privilege Level

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

Related Documentation

- Stacking a VLAN Tag on page 569
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

push-push

Syntax

push-push;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

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

Related Documentation

- Stacking Two VLAN Tags on page 570

premium (Output Priority Map)

Syntax

```
premium {
  forwarding-class class-name {
    loss-priority (high | low);
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Specifying an Output Priority Map on page 541
- input-priority-map on page 1200

premium (Policer)

Syntax

```
premium {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define a policer to apply to nonpremium traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Policers on page 539
- aggregate (Gigabit Ethernet CoS Policer) on page 1077
- ieee802.1p on page 1188

proxy

Syntax

proxy inet-address *address*;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family tcc],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.

Options

inet-address—Configure the IP address of the neighbor to the TCC router.

Required Privilege Level

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

Related Documentation

- Configuring Translation Cross-Connect Interface Switching on page 378
- *Example: Configuring an Ethernet TCC or Extended VLAN TCC*
- remote on page 1340
- *Junos OS VPNs Library for Routing Devices*

proxy-arp

Syntax proxy-arp (restricted | unrestricted);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
restricted added in Junos OS Release 10.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.

NOTE: You must configure the IP address and the inet family for the interface when you enable proxy ARP.

Default Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.

Options

- **none**—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
- **restricted**—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.
- **unrestricted**—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.

Default: **unrestricted**

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

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 369
- *Configuring Proxy ARP on Switches (CLI Procedure)*
- *Example: Configuring Proxy ARP on an EX Series Switch*
- Configuring Gratuitous ARP on page 19

push

Syntax push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

NOTE: On EX4300 switches, **push** is not supported at the [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map] hierarchy level.

Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.

You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

If you include the **push** statement in the configuration, you must also include the **pop** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level.

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

Related Documentation

- Stacking a VLAN Tag on page 569
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

push-push

Syntax push-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- Stacking Two VLAN Tags on page 570

quiet-period

Syntax quiet-period *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting authentication.

Options ***seconds***—Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting authentication.

Range: 0 through 65,535 seconds

Default: 60 seconds

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authenticator on page 1087**
- **dot1x on page 1124**
- **interface (IEEE 802.1x) on page 1205**

quiet-period (MX Series in Enhanced LAN Mode)

Syntax quiet-period *seconds*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*])]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X and captive portal authentication, configure the number of seconds the interface remains in the wait state following a failed authentication attempt by a supplicant before reattempting authentication.

Default 60 seconds

Options ***seconds***—Number of seconds the interface remains in the wait state.

Range: 0 through 65,535 seconds

Default: 60 seconds

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

reauthentication

Syntax reauthentication (disable | interval *seconds*);

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set or disable the periodic reauthentication of the client.

Options

- **disable**—Disable the periodic reauthentication of the client.
- **interval** ***seconds***—Specify the periodic reauthentication time interval.

Range: 1 through 65,535 seconds
Default: 3600 seconds

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- dot1x on page 1124
- interface (IEEE 802.1x) on page 1205
- quiet-period on page 1334

reauthentication (MX Series in Enhanced LAN Mode)

Syntax reauthentication *interval*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X authentication, specify the number of seconds before an authentication session times out.

Options *interval*—Sets the periodic reauthentication time interval in seconds.
Range: 1 through 4,294,967,296 seconds
Default: 3600 seconds

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

rebalance (Aggregated Ethernet Interfaces)

Syntax rebalance *interval*

Hierarchy Level [edit interfaces ae*X unit logical-unit-number* forwarding-options load-balance-stateful per-flow]

Release Information Statement introduced in Junos OS Release 13.2R1.

Description Configure periodic rebalancing of traffic flows of an aggregated Ethernet bundle by clearing the load balance state at a specified interval.

Options ***interval***—Number of minutes after which the load-balancing state must be cleared for the specified interface.

Range: 1 through 1000 flows per minute

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

Related Documentation

- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces on page 181

receive-options-packets

Syntax receive-options-packets;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

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

Related Documentation

- *Enabling Passive Monitoring on ATM Interfaces*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*

receive-ttl-exceeded

Syntax
receive-ttl-exceeded;

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

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

Related Documentation
- *Enabling Passive Monitoring on ATM Interfaces*
- *Enabling Packet Flow Monitoring on SONET/SDH Interfaces*

recovery

Syntax

```
recovery {
  (auto | manual);
  timer timer-value;
}
```

Hierarchy Level [edit interfaces *interfaces-name* link-degrade-monitor]

Release Information Statement introduced in Junos OS Release 15.1.

Description Configure the mechanism to be used to recover a degraded link. The recovery options supported are auto and manual.

Options **auto**—Recover a degraded link automatically. Use this option with the media-based action when there are no Layer 2 or Layer 3 protocols configured on the interface. If this option is configured, the degraded link is monitored at user-configured intervals; and if the link quality is found to have improved (if bit error rate hits the clear threshold), the link is automatically recovered. With this configuration, you must configure a timer value.

manual— Recover a degraded link manually. Use this option with the media-based action configuration when Layer 2 and Layer 3 protocols are configured on the interface. If this option is configured, you need to use the **request interface link-degrade-recover *interface-name*** statement to recover the link.

NOTE: The manual recovery option is recommended for user deployments that have static route configurations causing the remote end of the link to start forwarding packets (as soon as the physical link is up) while autorecovery is in progress.

timer *timer-value*—Specify the interval value (in seconds) after which autorecovery of the degraded link must be triggered. This option is applicable if you configure the autorecovery option. The interval period starts from the time the link is degraded. The default interval is 1800 seconds. The autorecovery attempt is repeated until the link is recovered or the link monitoring feature is disabled through configuration.

NOTE: During autorecovery, you might notice link flaps at the remote end of the link.

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

Related Documentation

- Link Degrade Monitoring Overview on page 381
- link-degrade-monitor on page 1221
- thresholds on page 1395
- request interface link-degrade-recover on page 1495

remote

Syntax

```
remote {
   (inet-address address | mac-address address);
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* family tcc],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the location of the remote router. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only.

Options

mac-address—Configure the MAC address of the remote site.

inet-address—Configure the IP address of the remote site.

Required Privilege Level

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

Related Documentation

- Configuring Translation Cross-Connect Interface Switching on page 378
- *Example: Configuring an Ethernet TCC or Extended VLAN TCC*
- proxy on page 1330
- *Junos OS VPNs Library for Routing Devices*

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 705

replay-window-size (MX Series)

Syntax replay-window-size *number-of-packets*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* replay-protect]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies the size of the replay protection window.

This statement has to be configured to enable replay protection.

When MACsec is enabled on an Ethernet link, an ID number is assigned to each packet entering the link. The ID number of the packet is checked by the receiving interface after the packet has traversed the MACsec-enabled link.

When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

Default Replay protection is disabled.

Options ***number-of-packets***—Specifies the size of the replay protection window, in packets.

When this variable is set to 0, all packets that arrive out-of-order are dropped.

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

replay-protect (MX Series)

Syntax

replay-protect {
  replay-window-size *number-of-packets*;
}

Hierarchy Level

[edit security macsec connectivity-association*connectivity-association-name*]

Release Information

Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description

Enable replay protection for MACsec.

A replay window size specified using the **replay-window-size** ***number-of-packets*** statement must be specified to enable replay protection.

Options

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 Media Access Control Security (MACsec) on MX Series Routers*

restore-interval

Syntax restore-interval *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Options *number*—Specify the restore interval.

Range: 1 through 12 minutes

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

retries

Syntax

retries *integer*;

Hierarchy Level

[edit protocols dot1x authenticator interface *interface-id*]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Set a limit on the number of failed authentication attempts between a port and a client. When the limit is exceeded, the port waits to reattempt authentication for the number of seconds set by the **quiet-period** statement configured at the same hierarchy level.

Options

integer—Specify the number of retries.

Range: 1 through 10

Default: 3 retries

Required Privilege Level

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **dot1x on page 1124**
- **interface (IEEE 802.1x) on page 1205**
- **quiet-period on page 1334**

retries (MX Series in Enhanced LAN Mode)

Syntax retries *number*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*])]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X and captive portal authentication, configure the number of times the switch attempts to authenticate the port after an initial failure. The port remains in a wait state during the quiet period after the authentication attempt.

Options *number*—Number of retries.

Default: 3 retries

Range: 1 through 10

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

revertive

Syntax revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Enable the ability to switch to a better priority link (if one is available).

NOTE: By default, LACP link protection is revertive. However, you can use this statement to define a specific aggregated Ethernet interface as revertive to override a global **non-revertive** statement specified at the **[edit chassis]** hierarchy level.

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

Related Documentation

- *non-revertive (Chassis)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

ring-id

Syntax ring-id *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify the ring ID.

Options *number*—Ring ID number.

Range: 1 through 239

Default: 1

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

ring-protection-link-end

Syntax ring-protection-link-end;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name* (east-interface | west-interface)]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

ring-protection-link-owner

Syntax ring-protection-link-owner;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*

routing-instance

Syntax

```
routing-instance {
  destination routing-instance-name;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.

Description To configure **interfaces** and **logical-systems**, specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default The default Internet routing table is **inet.0**.

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

Related Documentation

- *Junos OS Services Interfaces Library for Routing Devices*

routing-instance (PPPoE Service Name Tables)

Syntax

routing-instance *routing-instance-name*;

Hierarchy Level

[edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Use in conjunction with the **dynamic-profile** statement at the same hierarchy levels to specify the routing instance in which to instantiate a dynamic PPPoE interface. You can associate a routing instance with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The routing instance associated with a service entry in a PPPoE service name table overrides the routing instance associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **routing-instance** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **routing-instance** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options

routing-instance-name—Name of the routing instance in which the router instantiates the dynamic PPPoE interface.

Required Privilege Level

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

sa-multicast (100-Gigabit Ethernet)

Syntax sa-multicast;

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the 100-Gigabit Ethernet PIC or MIC to interoperate with other Juniper Networks 100-Gigabit Ethernet PICs.

NOTE: **The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.**

sa-multicast supports interoperability between the following PICs and MICs:

- 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP) and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) .
- 100-Gigabit Ethernet MICs and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4).

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

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 454
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 457
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode*
- *Interoperability Between MPC4E (MPC4E-3D-2CGE-8XGE) and 100-Gigabit Ethernet PICs on Type 4 FPC*
- *Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode*
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP on page 456
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 460
- **forwarding-mode (100-Gigabit Ethernet) on page 1174**

- sa-multicast (PTX Series Packet Transport Routers) on page 1351
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1444
- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 448

sa-multicast (PTX Series Packet Transport Routers)

Syntax sa-multicast;

Hierarchy Level [edit chassis fpc *slot* pic *slot* port *port-number* forwarding-mode]

Release Information Statement introduced in Junos OS Release 12.1X48R4.

Description Configure source address (SA) multicast bit mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP to enable interoperability with 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4.

> NOTE: When SA multicast bit steering mode is configured on a PTX Series Packet Transport Router 100-Gigabit Ethernet port, VLANs are not supported for that port.

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

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP on page 456
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 460

secure-authentication (MX Series in Enhanced LAN Mode)

Syntax secure-authentication (http | https);

Hierarchy Level [edit protocols captive-portal-custom-options]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Enable HTTP or HTTPS access on the captive portal interface.

Default **http**

Options **http**—Enables HTTP access on the captive portal interface.

https—Enables HTTPS access on the captive portal interface. HTTPS is recommended.

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

secure-channel

Syntax

```
secure-channel secure-channel-name {
  direction (inbound | outbound);
  encryption (MACsec);
  id {
    mac-address mac-address;
    port-idport-id-number;
  }
  offset (0|30|50);
  security-association security-association-number {
    key key-string;
  }
}
```

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name*]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Create and configure a secure channel to enable and configure MACsec when MACsec is enabled using static secure association key (SAK) security mode.

You do not need to use this option to enable MACsec using static connectivity association key (CAK) security mode. All configuration for MACsec using static CAK security mode is done inside of the connectivity association but outside of the secure channel. When MACsec is enabled using static CAK security mode, an inbound and an outbound secure channel—neither of which is user-configurable—is automatically created within the connectivity association.

Options 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 Media Access Control Security (MACsec) on MX Series Routers*

security-association

Syntax

security-association *security-association-number* {
  key *key-string*;
}

Hierarchy Level

[edit security macsec connectivity-association *connectivity-association-name* secure-channel *secure-channel-name*]

Release Information

Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description

Specifies the number of one of the security associations in the secure channel when MACsec is enabled using static secure association key (SAK) security mode. Because SAKs are created by the key server when MACsec is enabled using static connectivity association key (CAK) security mode, the **security-association** statement is not used when enabling MACsec using static CAK security mode.

You must configure at least two security associations to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.

Default

No security keys are configured, by default.

Options

security-association-number—Specifies the security association number and creates the SAK.

The security association number is a whole number between 0 and 3. You can configure two security associations in a secure channel when enabling MACsec using static security keys.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

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 700

server

Syntax server;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the router to operate in the PPPoE server mode. Supported on M120 and M320 Multiservice Edge Routers and MX Series Universal Edge Routers operating as access concentrators.

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

Related Documentation

- Configuring the PPPoE Server Mode on page 357

server-fail

Syntax server-fail (deny | permit | use-cache | *vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For EX Series switches configured for 802.1X authentication, specify the server fail fallback action the switch takes when all RADIUS authentication servers are unreachable.

When you specify the action ***vlan-name*** or ***vlan-id***, the VLAN must already be configured on the switch.

Default Authentication is denied.

Options **deny**—Force fail the supplicant authentication. No traffic will flow through the interface.

permit—Force succeed the supplicant authentication. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server.

use-cache—Force succeed the supplicant authentication only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected.

vlan-id—Move supplicant on the interface to the VLAN specified by this numeric identifier. This action is allowed only if it is the first supplicant connecting to the interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated.

vlan-name—Move supplicant on the interface to the VLAN specified by this name. This action is allowed only if it is the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated.

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

server-reject-vlan (MX Series in Enhanced LAN Mode)

Syntax

```
server-reject-vlan (vlan-id | vlan-name) {
   eapol-block;
   block-interval block-interval;
}
```

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For EX Series switches configured for 802.1X authentication, specify that when the switch receives an Extensible Authentication Protocol Over LAN (EAPoL) Access-Reject message during the authentication process between the switch and the RADIUS authentication server, supplicants attempting access to the LAN are granted access and moved to a specific VLAN. Any VLAN name or VLAN ID sent by a RADIUS server as part of the EAPoL Access-Reject message is ignored.

When you specify the VLAN ID or VLAN name, the VLAN must already be configured on the switch.

The remaining statements are explained separately. See CLI Explorer.

Default None

Options ***vlan-id***—Numeric identifier of the VLAN to which the supplicant is moved.

vlan-name—Name of the VLAN to which the supplicant is moved.

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

server-timeout

Syntax server-timeout *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Sets the number of seconds the port waits for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

Options ***seconds***—The number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

Range: 1 through 60 seconds

Default: 30 seconds

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authenticator on page 1087**
- **dot1x on page 1124**
- **interface (IEEE 802.1x) on page 1205**

server-timeout (MX Series in Enhanced LAN Mode)

Syntax server-timeout *seconds*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*])]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X and captive portal authentication, configure the amount of time a port will wait for a reply when relaying a response from the supplicant to the authentication server before timing out and invoking the server-fail action.

Default 30 seconds

Options *seconds* —Number of seconds.

Range: 1 through 60 seconds

Default: 30 seconds

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

service (PPPoE)

Syntax

```
service service-name {
  drop;
  delay seconds;
  terminate;
  dynamic-profile profile-name;
  routing-instance routing-instance-name;
  max-sessions number;
  agent-specifier {
    aci circuit-id-string ari remote-id-string {
      drop;
      delay seconds;
      terminate;
      dynamic-profile profile-name;
      routing-instance routing-instance-name;
      static-interface interface-name;
    }
  }
}
```

Hierarchy Level

[edit protocols pppoe service-name-tables *table-name*]

Release Information

Statement introduced in Junos OS Release 10.0.
any, **dynamic-profile**, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

Description

Specify the action taken by the interface on receipt of a PPPoE Active Discovery Initiation (PADI) control packet for the specified named service, **empty** service, or **any** service in a PPPoE service name table. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service.

Default

The default action is terminate.

Options

service-name—Service entry in the PPPoE service name table:

- *service-name*—Named service entry of up to 32 characters; for example, **premiumService**. You can configure a maximum of 512 named service entries across all PPPoE service name tables on the router.
- **empty**—Service entry of zero length that represents an unspecified service. Each PPPoE service name table includes one **empty** service entry by default.
- **any**—Default service for non-empty service entries that do not match the named or **empty** service entries configured in the PPPoE service name table. Each PPPoE service name table includes one **any** service entry by default.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag*
- *Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag*
- *Configuring the Action Taken for the Any Service*

service-name

Syntax service-name *name*;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description PPP over Ethernet interfaces, configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.

Options ***name***—Service to be requested from the PPP over Ethernet server.

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

Related Documentation

- Configuring the PPPoE Service Name on page 357
- *Junos OS Interfaces and Routing Configuration Guide*

service-name-table

Syntax

service-name-table *table-name*;

Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information

Statement introduced in Junos OS Release 10.0.
Support at the **[edit … family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description

Specify the PPPoE service name table assigned to a PPPoE underlying interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive.

NOTE: The **[edit … family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options

table-name—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

Required Privilege Level

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Assigning a Service Name Table to a PPPoE Underlying Interface*
- *Configuring the PPPoE Family for an Underlying Interface*

service-name-tables

Syntax

```
service-name-tables table-name {
  service service-name {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    max-sessions number;
    agent-specifier {
      aci circuit-id-string ari remote-id-string {
        drop;
        delay seconds;
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        static-interface interface-name;
      }
    }
  }
}
```

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 10.0.
dynamic-profile, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Create and configure a PPPoE service name table. Specify the action taken for each service and remote access concentrator on receipt of a PPPoE Active Discovery Initiation (PADI) packet. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service. A maximum of 32 PPPoE service name tables is supported per router.

Options ***table-name***—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Creating a Service Name Table*

session-expiry (MX Series in Enhanced LAN Mode)

Syntax session-expiry *seconds*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*])]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure the maximum duration in seconds of a session.

Options *seconds*—Duration of session.

Range: 1 through 65535

Default: 3600

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

source-address-filter

Syntax

```
source-address-filter {
  mac-address;
}
```

Hierarchy Level

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the **source-filtering** statement in the configuration to enable source address filtering.

Options

mac-address—MAC address filter. You can specify the MAC address as ***nn*:*nn*:*nn*:*nn*:*nn*:*nn*** or ***nnnn.nnnn.nnnn***, where ***n*** is a decimal digit. To specify more than one address, include multiple ***mac-address*** options in the **source-address-filter** statement.

If you enable the VRRP on a Fast Ethernet or Gigabit Ethernet interface, as described in "VRRP and VRRP for IPv6 Overview" on page 345, and if you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 3768, *Virtual Router Redundancy Protocol*. When you configure the VRRP group, the group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

Required Privilege Level

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

Related Documentation

- Configuring MAC Address Filtering for Ethernet Interfaces on page 14
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 16
- **source-filtering on page 1366**

source-filtering

Syntax (source-filtering | no-source-filtering);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the **source-address-filter** statement.

If the remote Ethernet card is changed, the interface is no longer able to receive packets from the new card because it has a different MAC address.

Default Source address filtering is disabled.

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

Related Documentation

- Configuring MAC Address Filtering for Ethernet Interfaces on page 14
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 16
- **accept-source-mac on page 1064**
- **source-address-filter on page 1365**

speed (Ethernet)

List of Syntax

Syntax (EX Series) on page 1367
Syntax (EX2300 and EX4300) on page 1367
Syntax (EX Series, ACX Series, MX Series) on page 1367
Syntax (QFX Series, OCX1100, EX4600) on page 1367

Syntax (EX Series)

speed (auto-negotiation | *speed*) ;

Syntax (EX2300 and EX4300)

speed *speed*;

Syntax (EX Series, ACX Series, MX Series)

speed (10m |10g | 100m | 1g | 2.5g |5g |auto | auto-10m-100m);

Syntax (QFX Series, OCX1100, EX4600)

speed (10g | 1g | 100m)

Hierarchy Level (EX Series)

[edit interfaces *interface-name* ether-options]

Hierarchy Level (EX2300 and EX4300)

[edit interfaces *interface-name*]

Hierarchy Level (ACX Series, EX Series, MX Series)

[edit interfaces *interface-name*],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Hierarchy Level (QFX Series, EX4600, OCX Series)

[edit interfaces *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Speed option 2.5Gbps introduced in Junos OS Release 18.1R2 for EX2300 switch.
Speed option 10Gbps and 5Gbps introduced in Junos OS Release 18.2R1 for EX4300 switch.

Description

Configure the interface speed. This statement applies to the management Ethernet interface (**fxp0** or **em0**), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, and Gigabit Ethernet interfaces on EX Series switches.

When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.

NOTE: On MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the **auto-10m-100m** option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the **show interfaces** command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the **show interfaces** command only. The actual interface speed is 100 Mbps.

On 10-Gigabit Ethernet SFP interfaces, autonegotiation is enabled by default and auto-detects the speed to be either 1 Gbps or 10 Gbps. On QFX5100-48S, QFX5100-96S, and QFX5100-24Q devices using 10-Gigabit Ethernet SFP interfaces, the speed is set to 10 Gbps by default and cannot be configured to operate in a different speed. On QFX5100-48S and QFX5100-96S devices using 1-Gigabit Ethernet SFP interfaces, the speed is set to 1 Gbps by default and cannot be configured to operate in a different speed.

NOTE: In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. You can, however, enable auto-negotiation by issuing the **set interface *name* ether-options auto-negotiation** command on the interface for which you want to change the interface speed. With autonegotiation enabled, the interface auto-detects the speed in which to operate.

NOTE: Only 10 Gbps and 40 Gbps interfaces are supported on OCX Series switches.

NOTE: When displaying interface information with **show interfaces** commands, you might see speed values for 1 Gbps interfaces displayed as **1000mbps**.

(For EX2300 only) Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches. The speed configuration statement is supported on both multi-rate gigabit ethernet interface (mge) and gigabit ethernet (ge) interface. The mge interface is a rate-selectable (multirate) Gigabit Ethernet interface that can support speeds of 10 Gbps, 5 Gbps, and 2.5 Gbps over CAT5e/CAT6/CAT6a cables. In the EX2300, the mge interface supports 100 Mbps, 1 Gbps, and 2.5 Gbps speeds, which can be configured by using the speed configuration statement. Note that 10Mbps speed is supported only on **ge** interfaces of EX2300 switch.

Default (EX Series) If the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled, the auto-negotiation option is enabled by default.

Options You can specify the speed as either **10m** (10 Mbps), **100m** (100 Mbps), and on MX Series routers, **1g** (1 Gbps). You can also specify the **auto** option on MX Series routers.

For Gigabit Ethernet interfaces on EX Series switches, you can specify one of the following options:

Table 100: Options for speed

Platforms	Speed Supported	Auto-negotiation
EX Series Switches	**100m**—100 Mbps **10m**—10 Mbps **1g**—1 Gbps	auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled.
ACX, MX Series	**100m**—100 Mbps **10m**—10 Mbps **1g**—1 Gbps	**auto**—Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link. **auto-10m-100m**—Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.
EX4600, QFX Series, QFabric, OCX100, QFX Series	**10g**—10 Gbps **1g**—1 Gbps **100m**—100 Mbps	auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled.
EX2300	**10m**—10 Mbps (supported on EX series switches and only on **ge** interfaces of EX2300 switch) **100m**—100 Mbps **1g**—1 Gbps **2.5g**—2.5 Gbps (supported only on **mge** interfaces of E2300 switch) **10g**—10 Gbps (supported only on **mge** interfaces for EX4300 switches) **5g**—5 Gbps (supported only on **mge** interfaces for EX4300 switches)	*speed*—Specify the interface speed. If the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the **auto-negotiation** statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end.

Required Privilege Level

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

Release History Table

Release	Description
18.2R1	Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches.

Related Documentation

- *Configuring the Interface Speed*
- Configuring the Interface Speed on Ethernet Interfaces on page 8
- Configuring Gigabit Ethernet Autonegotiation on page 551
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *auto-negotiation*
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX and QFX Series Switches*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

speed (MX Series DPC)

Syntax speed (auto | 1Gbps | 100Mbps | 10Mbps);

Hierarchy Level [edit interfaces ge-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 9.5.

Description On MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs you can set auto negotiation of speed. To specify the auto negotiation speed, use the **speed (auto | 1Gbps | 100Mbps | 10Mbps)** statement under the **[edit interface ge-/fpc/pic/port]** hierarchy level. The **auto** option will attempt to automatically match the rate of the connected interface. To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**.

NOTE: If the negotiated speed and the interface speed do not match, the link will not be brought up. Half duplex mode is not supported.

Options You can specify the speed as either **auto** (autonegotiate), **10Mbps** (10 Mbps), **100Mbps** (100 Mbps), or **1Gbps** (1 Gbps).

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

Related Documentation

- Configuring Gigabit Ethernet Autonegotiation on page 551
- **no-auto-mdix on page 1273**

stacked-vlan-tagging

Syntax stacked-vlan-tagging;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Description For Gigabit Ethernet IQ interfaces, Gigabit Ethernet, 10-Gigabit Ethernet LAN/WAN PIC, and 100-Gigabit Ethernet Type 5 PIC with CFP, enable stacked VLAN tagging for all logical interfaces on the physical interface.

For pseudowire subscriber interfaces, enable stacked VLAN tagging for logical interfaces on the pseudowire service.

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

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 559

static (Protocols 802.1X)

Syntax

```
static mac-address {
  interface interface-names;
  vlan-assignment (vlan-id |vlan-name );
}
```

Hierarchy Level [edit protocols authentication-access-control]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure MAC addresses to exclude from 802.1X authentication. The static MAC list provides an authentication bypass mechanism for supplicants connecting to a port, permitting devices such as printers that are not 802.1X-enabled to be connected to the network on 802.1X-enabled ports.

Using this 802.1X authentication-bypass mechanism, the supplicant connected to the MAC address is assumed to be successfully authenticated and the port is opened for it. No further authentication is done for the supplicant.

You can optionally configure the VLAN that the supplicant is moved to or the interfaces on which the MAC address can gain access from.

Options *mac-address* —The MAC address of the device for which 802.1X authentication should be bypassed and the device permitted access to the port.

The remaining statements are explained separately. See CLI Explorer.

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

static-interface

Syntax static-interface *interface-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Reserve the specified static PPPoE interface for use only by the PPPoE client with matching agent circuit identifier (ACI) and agent remote identifier (ARI) information. You can specify only one static interface per ACI/ARI pair configured for a named service entry, **empty** service entry, or **any** service entry in the PPPoE service name table.

The static interface associated with an ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.

If you include the **static-interface** statement in the configuration, you cannot also include either the **dynamic-profile** statement or the **routing-instance** statement. The **dynamic-profile**, **routing-instance**, and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options *interface-name*—Name of the static PPPoE interface reserved for use by the PPPoE client with matching ACI/ARI information. Specify the interface in the format **pp0.*logical***, where ***logical*** is a logical unit number from 0 through 16385 for static interfaces.

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

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client*

supplicant

Syntax supplicant *single*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the supplicant mode. Only single mode is supported.

This option will authenticate only the first client that connects to a port. All other clients that connect later (802.1x compliant or non-compliant) will be allowed free access on that port without any further authentication. If the first authenticated client logs out, all other users are locked out until a client authenticates again.

Options **single**—Sets single mode.

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authenticator on page 1087**
- **dot1x on page 1124**
- **interface (IEEE 802.1x) on page 1205**

supplicant (MX Series in Enhanced LAN Mode)

Syntax supplicant (multiple | single | single-secure);

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*])]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure the MAC-based method used to authenticate clients for 802.1X or captive portal authentication.

Default single

Options

single—Authenticates only the first client that connects to an authenticator port. All other clients connecting to the authenticator port after the first are permitted free access to the port without further authentication. If the first authenticated client logs out, all other supplicants are locked out until a client authenticates again.

single-secure—Authenticates only one client to connect to an authenticator port. The host must be directly connected to the switch.

multiple—Authenticates multiple clients individually on one authenticator port. You can configure the number of clients per port. If you also configure a maximum number of devices that can be connected to a port through port security settings, the lower of the configured values is used to determine the maximum number of clients allowed per port.

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

supplicant-timeout

Syntax supplicant-timeout *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

Options *seconds*—Specify the number of seconds the port waits for the supplicant timeout.

Range: 1 through 60 seconds

Default: 30 seconds

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- **authenticator on page 1087**
- **dot1x on page 1124**
- **interface (IEEE 802.1x) on page 1205**

supplicant-timeout (MX Series in Enhanced LAN Mode)

Syntax supplicant-timeout *seconds*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X authentication, configure how long the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request.

Default 30 seconds

Options ***seconds*** —Number of seconds.

Range: 1 through 60 seconds

Default: 30 seconds

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

swap

Syntax swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.

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

Related Documentation

- Rewriting the VLAN Tag on Tagged Frames on page 572
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

swap-by-poppush

Syntax swap-by-poppush;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 11.2

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to replace a VLAN tag. Pop original tag, then push an entirely new tag. The swap operation is performed as pop followed by push.

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

swap-push

Syntax swap-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with 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.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag on page 577

swap-swap

Syntax swap-swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and for 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- Rewriting the Inner and Outer VLAN Tags on page 578

switch-options

Syntax

```
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed (10m | 100m | 1g);
    link-mode (full-duplex | half-duplex);
  }
}
```

Hierarchy Level [edit interfaces ge-*pim*/0/0]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configuration of the physical port characteristics is done under the single physical interface.

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

switch-port

Syntax

```
switch-port port-number {
  (auto-negotiation | no-auto-negotiation);
  speed (10m | 100m | 1g);
  link-mode (full-duplex | half-duplex);
}
```

Hierarchy Level [edit interfaces ge-*pim*/0/0 switch-options]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configuration of the physical port characteristics, done under the single physical interface.

Default Autonegotiation is enabled by default. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation.

Options ***port-number***—Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

The remaining statements are explained separately. See CLI Explorer.

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

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 694
- Configuring Threshold Values for Fault Events in an Action Profile on page 703

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 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

Description

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

Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Specifying the Actions to Be Taken for Link-Fault Management Events on page 700
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*

system-id

Syntax system-id *system-id*;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp]

Release Information Statement introduced in Junos OS Release 12.2R1
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define the LACP system identifier at the aggregated Ethernet interface level.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

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

Related Documentation

- Configuring LACP for Aggregated Ethernet Interfaces on page 140

system-priority

Syntax system-priority *priority*;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global **[edit chassis]** hierarchy level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options ***priority***—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

Range: 0 through 65535

Default: 127

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

tag-protocol-id (TPIDs Expected to Be Sent or Received)

Syntax

tag-protocol-id [*tpids*];

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile],
[edit interfaces *interface-name* ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the **tag-protocol-id** statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (**0x8100**) is supported.

For Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, define the TPIDs expected to be sent or received on a particular VLAN. The default TPID value is **0x8100**. Other supported values are **0x88a8**, **0x9100**, and **0x9200**.

Options

tpids—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.

Required Privilege Level

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

Related Documentation

- Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames on page 563
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*

tag-protocol-id (TPID to Rewrite)

Syntax tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers the default TPID value (**0x8100**) is supported.

Default If the **tag-protocol-id** statement is not configured, the TPID value is 0x8100.

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

Related Documentation

- Configuring Inner and Outer TPIDs and VLAN IDs on page 566

targeted-options (Grouping Subscribers by Bandwidth Usage)

Syntax

```
targeted-options {
  backup backup;
  group group;
  primary primary;
  weight weight-value;
}
```

Hierarchy Level

[edit dynamic-profiles *name* interfaces *name* unit],
[edit dynamic-profiles *name* logical-systems *name* interfaces *name* unit],
[edit interfaces *name* unit *unit-number*]

Release Information

Statement introduced in Junos OS Release 16.1.
weight statement introduced in Junos OS Release 17.3 for MX240, MX480, MX960, and MX Virtual Chassis.

Description

Configure primary and backup links for manual targeting.

Options

backup—Specify a backup member link per subscriber when you configure manual targeting. Configuring a backup link is optional.

group—(Optional) Assign a group name for subscribers with similar bandwidth usage. Subscribers that are configured for targeted distribution without a group name are added to the **default** group and distributed evenly across member links. Grouping of subscribers is supported only for static subscribers.

Default: default

primary—Specify a primary member link per subscriber when you configure manual targeting. You must always configure a primary link when you configure manual targeting.

weight ***weight-value***—Specify the weight for all the targeted subscribers like PPPoe, demux, and conventional VLANs based on the customers preferences, class of service (CoS), or bandwidth requirement. Member links for logical interfaces of aggregated Ethernet logical interfaces are assigned based on the value of the weight . When a new VLAN is added to the same aggregated Ethernet bundle, then the primary member link selected for targeting is the one with the minimum primary load and the backup link selected for targeting is the one with the minimum overall load.

NOTE: **The weight *weight-value* statement is supported only at the [edit interfaces *name* unit *unit-number*] hierarchy level.**

Required Privilege Level

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

Related Documentation

- *Understanding Support for Targeted Distribution of Logical Interface Sets of Static VLANs over Aggregated Ethernet Logical Interfaces*
- targeted-options (Manual Targeting) on page 1391
- *weight*

targeted-options (Manual Targeting)

Syntax

```
targeted-options {
  (logical-interface-chassis-redundancy | logical-interface-fpc-redundancy);
  rebalance-periodic {
    interval interval;
    start-time start-time;
  }
  type (auto | manual);
}
```

Hierarchy Level

[edit dynamic-profiles *name* interfaces *name* aggregated-ether-options],
[edit dynamic-profiles *name* logical-systems *name* interfaces *name*
aggregated-ether-options],
[edit interfaces *name* aggregated-ether-options]
[edit interfaces *name* unit]

Release Information Statement introduced in Junos OS Release 16.1.

Description Configure manual targeting or auto-targeting.

Options

type—Configure manual targeting type as **manual** or **auto**.

Values:

- **auto**—Configure targeted-distribution without specific primary and backup links.
- **manual**—Configure targeted distribution with specific member links as primary and backup for a subscriber. When you configure manual targeting, you must always configure a primary link. Configuring a backup link is optional. You specify the primary and backup links for a subscriber in the individual interface configuration. You configure primary and backup links by using the targeted-options statement at the **[edit interfaces *name* unit]** hierarchy level.

  Manual targeting enhances the distribution of targeted VLANs or subscribers across member links of an aggregated Ethernet bundle by making it bandwidth-aware.

Default: **auto**

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

- Ethernet Interfaces Overview on page 3
- targeted-options (Grouping Subscribers by Bandwidth Usage) on page 1389
- *Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis*

targeted-distribution

Syntax targeted-distribution primary-list *primary-list* |backup-list *backup-list;*

Hierarchy Level [edit logical-systems *name* interfaces *name* unit]

Release Information Statement introduced in Junos OS Release 16.1R1.

Description Configure egress data for a member link in an aggregated Ethernet bundle. Specify a distribution list as primary list and a different distribution list as backup list. A backup list is provisioned in the event the primary list goes down.

Options **primary-list**—(Optional) Specify the role of the distribution list as primary. Member links of the aggregated Ethernet are assigned membership to the distribution list.

backup-list—(Optional) Specify the role of the distribution list as backup. Member links of the aggregated Ethernet are assigned membership to the distribution list.

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

Related Documentation

- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links on page 202
- **distribution-list on page 1122**
- **targeted-options on page 1393**

targeted-options

Syntax

```
targeted-options {
  type (auto | manual);
}
```

Hierarchy Level

[edit dynamic-profiles *name* interfaces *name* aggregated-ether-options]
[edit dynamic-profiles *name* logical-systems *name* interfaces *name* aggregated-ether-options]
[edit interfaces *name* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Specify the type of targeting to be used for targeted distribution. Specify the targeting option as **manual** for conventional VLAN targeting. By default, the targeting option is **auto**.

Options

type—Specify the type of targeting to be used for targeted distribution.

Default: **auto**—By default, targeted option is set to **auto**.

Values:

- **manual**—Use **manual** keyword to enforce manual targeting on conventional VLANs.

Required Privilege Level

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

Related Documentation

- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links on page 202
- **distribution-list on page 1122**
- **targeted-distribution on page 1392**

terminate (PPPoE Service Name Tables)

Syntax

terminate;

Hierarchy Level

[edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information

Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description

Direct the router to immediately respond to a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client by sending the client a PPPoE Active Discovery Offer (PADO) packet. The PADO packet contains the name of the access concentrator (router) that can service the client request. The **terminate** action is the default action for a named service entry, **empty** service entry, **any** service entry, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in a PPPoE service name table.

Required Privilege Level

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

Related Documentation

- *Configuring PPPoE Service Name Tables*

thresholds

Syntax

```
thresholds {
  clear clear-value;
  interval interval-value;
  set set-value;
  warning-clear warning-clear-value;
  warning-set warning-set-value;
}
```

Hierarchy Level [edit interfaces *interfaces-name* link-degrade-monitor]

Release Information Statement introduced in Junos OS Release 15.1.

Description Configure the BER threshold values (such as set and clear thresholds) at which different corrective actions must be triggered on a degraded interface.

Options

clear *clear- value*—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 12.

interval *interval-value*—The number of consecutive link degrade events that are considered before any corrective action is taken. The supported value range for the interval is 1 through 256, and the default interval is 10.

set *set- value*—The BER threshold value at which the link is considered degraded and a corrective action, specified by the user, is triggered. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 7.

warning clear *warning-clear-value*—The link clear warning threshold. Every time this threshold value is reached, a system message is logged to indicate that the link degrade condition has been cleared on the interface. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 11.

warning set *warning-set-value*—The link degrade warning threshold. Every time this threshold value is reached, a system message is logged to indicate that a link degrade has occurred on the interface. You can configure this value in the 1E–*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of IE–3 refers to the BER threshold value of $1x10^{-3}$. The supported exponent range is 1 through 16, and the default value is 9.

NOTE: The lower the BER with high confidence level, the longer it takes to estimate it. In such cases, a few packet drops might be noticed (based on the bit error distribution) before a link degrade event is detected.

Required Privilege Level

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

Related Documentation

- Link Degrade Monitoring Overview on page 381
- link-degrade-monitor on page 1221
- recovery on page 1339
- request interface link-degrade-recover on page 1495

traceoptions

Syntax

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

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description Set LLDP protocol-level tracing options.

Default The default LLDP protocol-level trace options are inherited from the global **traceoptions** statement.

Options **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place spanning-tree protocol tracing output in the file **/var/log/stp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the LLDP-specific tracing options:

- **all**—Trace all operations.
- **config**—Log configuration events.
- **interface**—Trace interface update events.
- **protocol**—Trace protocol information.
- **rtsock**—Trace socket events.

- **vlan**—Trace vlan update events.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events. This is the default. If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file. This is the default. If you do not include this option, tracing output is appended to an existing trace file.

size *maximum-file-size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

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

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

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

Related Documentation

- Tracing LLDP Operations on page 344

traceoptions (Individual Interfaces)

List of Syntax

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) on page 1399
Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) on page 1399
Syntax (OAMLFM with EX Series, QFX Series, NFX Series) on page 1399
Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) on page 1399

Syntax (Individual interfaces with PTX Series, EX Series, ACX Series)

```
traceoptions {
  file filename <files name> <size size> <world-readable | no-world-readable>;
  flag flag;
  match;
}
```

Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series)

```
traceoptions {
  flag flag;
}
```

Syntax (OAMLFM with EX Series, QFX Series, NFX Series)

```
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
```

Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series)

[edit interfaces *interface-name*]

Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

[edit interfaces]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in JUNOS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple **flag** statements.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** file in the directory **/var/log/dcd**.

On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.

On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).

Default If you do not include this statement, no interface-specific tracing operations are performed.

Table 101: OptiOptionsfor traceoptions

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
file*filename*	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log/dcd**. By default, interface process tracing output is placed in the file.	-	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.	—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file **dcd**
files *number*	**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. Then the oldest trace file is overwritten.	-	—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.	—(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. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option. *Range:* 2 through 1000 *Default:* 3 files
flag	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options. • **all**—All interface tracing operations • **event**—Interface events • **ipc**—Interface interprocess	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options. • **all**—All interface tracing operations • **event**—Interface events • **ipc**—Interface interprocess	—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: • action-profile—Trace action profile invocation events. • all—Trace all events. • configuration—Trace configuration events. • protocol—Trace	—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags: • **all** • **change-events**—Log changes that produce configuration events • **config-states**—Log the configuration

Table 101: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
	communication (IPC) messages • **media**—Interface media changes • **q921**—Trace ISDN Q.921 frames • **q931**—Trace ISDN Q.931 frames	communication (IPC) messages • **media**—Interface media changes • **q921**—Trace ISDN Q.921 frames • **q931**—Trace ISDN Q.931 frames	protocol processing events. • routing socket—Trace routing socket events.	state machine changes • **kernel**—Log configuration IPC messages to kernel • **kernel-detail**—Log details of configuration messages to kernel
match	—(Optional) Regular expression for lines to be traced.	-	—(Optional) Refine the output to log only those lines that match the given regular expression.	-
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 this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.	-	—(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 GB *Range:* 10 KB through 1 GB *Default:* 128 KB *Default:* If you do not include this option, tracing output is appended to an existing trace file.	

Table 101: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
				—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option. *Syntax:* ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes *Range:* 10 KB through the maximum file size supported on your router *Default:* 1 MB
no-world-readable	—(Optional) Prevent any user from reading the log file.	-	—(Optional) Restrict file access to the user who created the file.	—(Optional) Disallow any user to read the log file.
world-readable	—(Optional) Allow any user to read the log file.	-	—(Optional) Enable unrestricted file access.	—(Optional) Allow any user to read the log file.
disable	-	-	-	

Table 101: Options for traceoptions (continued)

Option	Individual interfaces with PTX Series, ACX Series, EX Series	Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series	Interface Process with OAMLFM with EX Series, QFX Series, NFX Series	Interface process with ACX Series, SRX Series, MX Series, M Series, T Series
				—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.
no-remote-trace	-	-	—(Optional) Disable the remote trace.	-
match ***regex***	-	-	-	—(Optional) Refine the output to include only those lines that match the given regular expression.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Tracing Operations of an Individual Router Interface*
- *Tracing Operations of an Individual Router or Switch Interface*
- *Example: Configuring Ethernet OAM Link Fault Management*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*
- *Tracing Operations of the Interface Process*

traceoptions (LACP)

Syntax

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

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Description Define tracing operations for the LACP protocol.

Default If you do not include this statement, no LACP protocol tracing operations are performed.

Options ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, interface process tracing output is placed in the file **lacpd**.

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. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

no-world-readable—(Optional) Prevent any user from reading the log 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 this size, it is renamed ***trace-file*.0.** When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option:

Syntax: ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

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

Related Documentation

- Tracing LACP Operations on page 146

traceoptions (MACsec)

Syntax

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

Hierarchy Level

[edit security macsec]

Release Information

Statement introduced in Junos OS Release 15.1 for MIC-3D-20GE-SFP-E on MX Series routers.
Statement introduced in Junos OS Release 16.1 for MPC7E-10G on MX Series routers.
Statement introduced in Junos OS Release 17.3R2 for JNP-MIC1-MACSEC MIC on MX10003 routers.

Description

Define tracing operations at the MACsec level. Tracing operations provide support for debugging protocol-level issues. MACsec is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. To specify more than one tracing operation, include multiple **flag** statements.

The interfaces **traceoptions** statement does not support a separate trace file. The logging is done by the kernel, so the tracing information is placed in the **syslog** file in the directory **/var/log/dcd**.

Default

If you do not include this statement, no tracing operations are performed.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. By default, interface process tracing output is placed in the directory. If you do not specify the name of the trace file, all files are placed in the directory **/var/log/dcd**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches the maximum value, 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. Values range from 2 through 1000.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the tracing operation options:

all—Trace all operations.

config—Trace configuration messages.

debug—Trace debug messages.

normal—Trace normal messages.

no-world-readable—(Optional) Prevent any user from reading the log 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 this size, it is renamed **trace-file.0**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Understanding Media Access Control Security (MACsec) on MX Series Routers*
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

traceoptions (MACsec interfaces)

Syntax

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

Hierarchy Level [edit security macsec interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 16.1 for MPC7E-10G on MX Series routers. Statement introduced in Junos OS Release 17.3R2 for JNP-MIC1-MACSEC MIC on MX10003 routers.

Description Define tracing operations for individual MACsec interfaces. Tracing operations provide support for debugging protocol-level issues. MACsec is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. To specify more than one tracing operation, include multiple **flag** statements.

The interfaces **traceoptions** statement does not support a separate trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** file in the directory **/var/log/dcd**.

Default If you do not include this statement, no tracing operations are performed.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. By default, interface process tracing output is placed in the directory. If you do not specify the name of the tracefile, all files are placed in the directory **/var/log/dcd**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named trace-file reaches the maximum value, 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. Values range from 2 through 1000.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the tracing operation options:

all—Trace all operations.

keys—Trace key creation or generation information.

mka-packets—Trace MACsec Key Agreement (MKA) protocol input and output packet information.

normal—Trace all normal events and messages.

state—Trace MKA protocol state information.

to-secy—Trace MKA to security entity state change information.

no-world-readable—(Optional) Prevent any user from reading the log 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 this size, it is renamed **trace-file.0**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- *Understanding Media Access Control Security (MACsec) on MX Series Routers*
- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

traceoptions (PPPoE)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    aci regular-expression;
    ari regular-expresion;
    service-name regular-expresion;
    underlying-interface interface-name;
  }
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 9.6.
Option **filter** introduced in Junos OS Release 12.3

Description Define tracing operations for PPPoE processes.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

disable—Disable this trace flag.

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

BEST PRACTICE: Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.

- **aci *regular-expression***—Regular expression to match the agent circuit identifier provided by PPPoE client.

- **ari** ***regular-expression***—Regular expression to match the agent remote identifier provided by PPPoE client.
- **service** ***regular-expression***—Regular expression to match the name of PPPoE service.
- **underlying-interface** ***interface-name***—Name of a PPPoE underlying interface. You cannot use a regular expression for this filter option.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **config**—Trace configuration events.
- **events**—Trace events.
- **gres**—Trace GRES events.
- **init**—Trace initialization events.
- **interface-db**—Trace interface database operations.
- **memory**—Trace memory processing events.
- **protocol**—Trace protocol events.
- **rtsock**—Trace routing socket events.
- **session-db**—Trace connection events and flow.
- **signal**—Trace signal operations.
- **state**—Trace state handling events.
- **timer**—Trace timer processing.
- **ui**—Trace user interface processing.

level—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

Default: **error**

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- Tracing PPPoE Operations on page 360

traceoptions (802.1X and Captive Portal for MX Series in Enhanced LAN Mode)

Syntax

traceoptions {
  file *filename* <files *number*> <size *size*> <world-readable | no-world-readable> <match
    regex>;
  flag *flag*;
}

Hierarchy Level

[edit protocols authentication-access-control]

Release Information

Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description

Define tracing operations for the 802.1X protocol, LLDP, and captive portal authentication.

Default

Tracing operations are disabled.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum 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 by using 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**—All tracing operations.
- **config-internal**—Trace internal configuration operations.
- **dot1x-event**—Trace 802.1x events.
- **dot1x-debug**—Trace 802.1x events.
- **dot1x-ipc**—Trace IPC interactions.
- **eapol**—Trace EAPOL packets transmitted and received.
- **general**—Trace general operations.
- **normal**—Trace normal operations.
- **parse**—Trace reading of the configuration.
- **regex-parse**—Trace regular-expression parsing operations.
- **state**—Trace protocol state changes.

- **task**—Trace protocol task operations.
- **timer**—Trace protocol timer operations.
- **vlan**—Trace VLAN transactions.

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Restrict 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 with the **files** option, you also must specify a maximum file size.

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

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

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

transmit-delay

Syntax

transmit-delay *seconds*;

Hierarchy Level

[edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

(MX Series and T Series routers only) Configure a delay between two successive LLDP advertisements.

Options

seconds—Delay between two successive LLDP advertisements.

Default: 2

Range: 1 through 8192

Required Privilege Level

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

Related Documentation

- Configuring LLDP on page 338

transmit-interval (MACsec for MX Series)

Syntax transmit-interval *interval*;

Hierarchy Level [edit security macsec connectivity-association *connectivity-association-name* mka]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description Specifies the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs).

The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower *interval* increases bandwidth overhead on the link; a higher *interval* optimizes the MKA protocol data unit exchange process.

The transmit interval settings must be identical on both ends of the link when MACsec using static connectivity association key (CAK) security mode is enabled.

We recommend increasing the interval to 6000 ms in high-traffic load environments.

Default The default transmit interval is 2000 milliseconds.

Options ***interval***—Specifies the transmit interval, in milliseconds.

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

Related Documentation

- *Configuring Media Access Control Security (MACsec) on MX Series Routers*

transmit-period

Syntax transmit-period *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in Junos OS Release 9.3.

Description Set the number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Options *seconds*—The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Range: 1 through 65,535 seconds

Default: 30 seconds

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

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 33
- authenticator on page 1087
- dot1x on page 1124
- interface (IEEE 802.1x) on page 1205

transmit-period (MX Series in Enhanced LAN Mode)

Syntax transmit-period *seconds*;

Hierarchy Level [edit protocols authentication-access-control interface (all | [*interface-names*]) dot1x]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description For 802.1X authentication, how long the port waits before retransmitting the initial EAPOL PDUs to the supplicant.

Default 30 seconds

Options *seconds*—Number of seconds the port waits before retransmitting the initial EAPOL PDUs to the supplicant.

Range: 1 through 65,535 seconds

Default: 30 seconds

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

uac-policy (MX Series in Enhanced LAN Mode)

Syntax uac-policy;

Hierarchy Level [edit protocols authentication-access-control]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure Junos Pulse Access Control Service as the access policy to authenticate and authorize users connected to the switch for admission to the network and for access to protected network resources.

Default The Access Control Service access policy is disabled.

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

Related Documentation

underlying-interface

Syntax underlying-interface *interface-name*;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Release Information Statement introduced before Junos OS Release 7.4.
Support for aggregated Ethernet added in Junos OS Release 9.4.

Description Configure the interface on which PPP over Ethernet is running.

For demux interfaces, configure the underlying interface on which the demultiplexing (demux) interface is running.

Options *interface-name*—Name of the interface on which PPP over Ethernet or demux is running. For example, **at-0/0/1.0** (ATM VC), **fe-1/0/1.0** (Fast Ethernet interface), **ge-2/0/0.0** (Gigabit Ethernet interface), **ae1.0** (for IP demux on an aggregated Ethernet interface), or **ae1** (for VLAN demux on an aggregated Ethernet interface).

NOTE: **Demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet interfaces, or aggregated Ethernet devices.**

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

Related Documentation

- *Configuring an IP Demultiplexing Interface*
- *Configuring a VLAN Demultiplexing Interface*
- Configuring the PPPoE Underlying Interface on page 356
- *Junos OS Interfaces and Routing Configuration Guide*

unit

Syntax

```
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  auto-configure {
    agent-circuit-identifier {
      dynamic-profile profile-name;
    }
    line-identity {
      include {
        accept-no-ids;
        circuit-id;
        remote-id;
      }
      dynamic-profile profile-name;
    }
  }
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      maximum-contexts number <force>;
      f-max-period number;
      queues [queue-numbers];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
  compression-device interface-name;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  demux-source family;
  demux-options {
    underlying-interface interface-name;
```



```
}
description text;
etree-ac-role (leaf | root);
interface {
  l2tp-interface-id name;
  (dedicated | shared);
}
dialer-options {
  activation-delay seconds;
  callback;
  callback-wait-period time;
  deactivation-delay seconds;
  dial-string [dial-string-numbers];
  idle-timeout seconds;
  incoming-map {
    caller caller-id | accept-all;
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
      [routes];
    }
  }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
family family-name {
  ... the family subhierarchy appears after the main [edit interfaces interface-name unit
    logical-unit-number] hierarchy ...
}
fragment-threshold bytes;
host-prefix-only;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap |
  swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
```

```
    output-policer policer-name;
    output-three-color policer-name;
  }
  link-layer-overhead percent;
  minimum-links number;
  mrru bytes;
  multicast-dlci dlci-identifier;
  multicast-vci vpi-identifier.vci-identifier;
  multilink-max-classes number;
  multipoint;
  oam-liveness {
    up-count cells;
    down-count cells;
  }
  oam-period (disable | seconds);
  output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap |
    swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
  }
  passive-monitor-mode;
  peer-unit unit-number;
  plp-to-clp;
  point-to-point;
  ppp-options {
    mru size;
    mtu (size | use-lower-layer);
    chap {
      access-profile name;
      default-chap-secret name;
      local-name name;
      passive;
    }
    compression {
      acfc;
      pfc;
    }
    dynamic-profile profile-name;
    ipcp-suggest-dns-option;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-restart-timer milliseconds;
    pap {
      access-profile name;
      default-pap-password password;
      local-name name;
      local-password password;
      passive;
    }
  }
  pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
```



```
  service-name name;
  underlying-interface interface-name;
}
pppoe-underlying-options {
  access-concentrator name;
  direct-connect;
  dynamic-profile profile-name;
  max-sessions number;
}
proxy-arp;
service-domain (inside | outside);
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
    length);
  queue-length number;
}
short-sequence;
targeted-distribution;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
  backup-destination address;
  destination address;
  key number;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      (input | output | input output);
    }
  }
  access-concentrator name;
  address address {
    ... the address subhierarchy appears after the main [edit interfaces interface-name unit
      logical-unit-number family family-name] hierarchy ...
  }
  bundle interface-name;
  core-facing;
  demux-destination {
    destination-prefix;
  }
  demux-source {
```

```
    source-prefix;
}
direct-connect;
duplicate-protection;
dynamic-profile profile-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list [filter-names];
    output filter-name;
    output-list [filter-names];
}
interface-mode (access | trunk);
ipsec-sa sa-name;
keep-address-and-control;
mac-validate (loose | strict);
max-sessions number;
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name
    mode loose;
}
sampling {
    input;
    output;
}
service {
    input {
        post-service-filter filter-name;
        service-set service-set-name <service-filter filter-name>;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
service-name-table table-name
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address
    destination-profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
```



```
        address address {
          arp ip-address (mac | multicast-mac) mac-address <publish>;
          broadcast address;
          destination address;
          destination-profile name;
          eui-64;
          master-only;
          multipoint-destination address {
            dlci dlci-identifier;
            epd-threshold cells <plp1 cells>;
            inverse-arp;
            oam-liveness {
              up-count cells;
              down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
              (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
                sustained rate);
              queue-length number;
            }
            vci vpi-identifier.vci-identifier;
          }
          preferred;
          primary;
          (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise–interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
              hold-time seconds;
            }
            priority number;
            track {
              interface interface-name {
                bandwidth-threshold bits-per-second priority-cost number;
              }
              priority-hold-time seconds;
              route ip-address/prefix-length routing-instance instance-name priority-cost cost;
            }
            virtual-address [addresses];
            virtual-link-local-address ipv6–address;
            vrrp-inherit-from {
              active-interface interface-name;
              active-group group-number;
            }
          }
        }
      }
    }
```

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*],
[edit interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number***—Number of the logical unit.

Range: 0 through 1,073,741,823 for demux and PPPoE static interfaces. 0 through 16,385 for all other static interface types.

etree-ac-role (leaf | root)—To configure an interface as either leaf or root.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring Logical Interface Properties*
- *Junos OS Services Interfaces Library for Routing Devices*

unnumbered-address (Dynamic Profiles)

Syntax unnumbered-address *interface-name* <preferred-source-address *address*>;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

Release Information Statement introduced in Junos OS Release 9.2.
Support for the **$junos-preferred-source-address** and **$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.
Support for the **$junos-loopback-interface** predefined variable introduced in Junos OS Release 9.6.

Description For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the **$junos-loopback-interface-address** predefined variable.

You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.

Options *interface-name*—Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the **$junos-loopback-interface** predefined variable.

When defining the **unnumbered-address** statement using a static interface, keep the following in mind:

- If you choose to include the **routing-instance** statement at the **[edit dynamic-profiles]** hierarchy level, that statement must be configured with a dynamic value by using the **$junos-routing-instance** predefined variable. In addition, whatever static unnumbered interface you specify must belong to that routing instance; otherwise, the profile instantiation fails.
- If you choose to not include the **routing-instance** statement at the **[edit dynamic-profiles]** hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance.

NOTE: When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the $junos-routing-instance predefined variable, you must not configure a preferred source address, whether with the $junos-preferred-source-address predefined variable, the $junos-preferred-source-ipv6-address predefined variable, or the

preferred-source-address statement. Configuring the preferred source address in this circumstance causes a commit failure.

When defining the **unnumbered-address** statement using the **$junos-loopback-interface** predefined variable, keep the following in mind:

- To use the **$junos-loopback-interface** predefined variable, the dynamic profile must also contain the **routing-instance** statement configured with the **$junos-routing-instance** predefined variable at the [edit dynamic-profiles] hierarchy level.
- The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

address—(Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the **$junos-preferred-source-address** predefined variable for the inet family, or the **$junos-preferred-source-ipv6-address** predefined variable for the inet6 family.

When defining the **preferred-source-address** value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the **preferred-source-address** must be configured in the specified unnumbered interface.

When defining the **preferred-source-address** value using the **$junos-preferred-source-address** or the **$junos-preferred-source-ipv6-address** predefined variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **$junos-loopback-interface** predefined variable.
- You must configure the **routing-instance** statement using the **$junos-routing-instance** predefined variable at the **[edit dynamic-profiles]** hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

Required Privilege Level

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

Related Documentation

- *Dynamic Profiles Overview*

unnumbered-address (PPP)

Syntax unnumbered-address *interface-name* destination *address* destination-profile *profile-name*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

Release Information Statement introduced before Junos OS Release 7.4.

Description For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.

Options ***interface-name***—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details.

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

Related Documentation

- *Configuring IPCP Options for Interfaces with PPP Encapsulation*
- *Junos OS Administration Library*

version-3

Syntax version-3;

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 12.2.

Description Enable Virtual Router Redundancy Protocol version 3 (VRRPv3).

NOTE:

- Even though the **version-3** statement can be configured only at the **[edit protocols vrrp]** hierarchy level, VRRPv3 is enabled on all the configured logical systems as well.
- When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.

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

Related Documentation
- *Junos OS Support for VRRPv3*

virtual-control-channel

Syntax

```
virtual-control-channel channel-name {
  west-interface name;
  east-interface name;
}
```

Hierarchy Level [edit protocols protection-group ethernet-ring *name* (east-interface | west-interface)]

Release Information Statement introduced in Junos OS Release 14.2.

Description Specify virtual control channels which are logical interfaces on the east and west interfaces of the major ring.

Options **west-interface** ***name***—Logical interface on the major ring's west port.

east-interface ***name***—Logical interface on the major ring's east port.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

virtual-switch

Syntax virtual-switch *name* bridge-domain *name* vlan-id [*vlan-ids*];

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* default-*x*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the routing-instance type as a virtual switch, under which bridge-domain MIPs must be enabled.

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

Related Documentation

- *Configuring MIP for Bridge Domains of a Virtual Switch*

vlan-assignment

Syntax vlan-assignment (*vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols authentication-access-control]

Release Information Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Configure the VLAN that is associated with the list of MAC addresses that are excluded from RADIUS authentication.

Options ***vlan-id*** | ***vlan-name***—The name of the VLAN or the VLAN tag identifier to associate with the device. The VLAN already exists on the switch.

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

vlan-id (VLAN ID to Be Bound to a Logical Interface)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options ***number***—A valid VLAN identifier.

Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.

For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.

VLAN ID 0 is reserved for tagging the priority of frames.

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

Related Documentation

- Enabling VLAN Tagging on page 247

vlan-id

Syntax

vlan-id (all | none | *number*);

VLANs and Bridge Domain VLANs

For platforms without ELS:

[edit vlans *vlan-name* vlan-range]

For platforms without ELS and with ELS:

[edit vlans *vlan-name*]

For ELS platforms only:

[edit interfaces *interface-name* unit *number*]
[edit vlans *vlan-name* vlan-id-list]

[edit vlans *vlan-name*],
[edit logical-systems *logical-system-name* vlans*vlan-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* vlans*vlan-name*],
[edit routing-instances *routing-instance-name* vlans *vlan--name*]

802.1Q Tagging

[edit vlans *vlan-name*]

VLAN ID to Rewrite

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

VLAN Tagging and Layer 3 Subinterfaces

[edit interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.2 for EX Series switches VLAN tagging and Layer 3 subinterfaces.
Support for Layer 2 trunk ports added in Junos OS Release 9.2.
Support for SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For VLANs, specify a VLAN identifier (VID) to include in the packets sent to and from the VLAN, or a VPLS routing instance.

NOTE: When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the none option are not permitted.

For 802.1Q tagging, configure an 802.1Q tag to apply to all traffic that originates on the VLAN.

The number zero is reserved for priority tagging and the number 4095 is also reserved.

For VLAN ID to Rewrite Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the **vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **vlan-id** statement that you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Default For 802.1Q Tagging on EX Series and SRX Series, If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1.

For VLANs on a QFX3500 and QFX3500 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.

On a QFX5100 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.

NOTE: You can only create up 4090 VLANs on a QFX5100 switch. If you create more than 4090 VLANs, the interfaces associated with the extra VLANs are not displayed in the **show vlans** command output. For example, if you create 4094 VLANs, the extra VLANs will not have interfaces associated with the VLANs. The order in which you configure the extra VLANs determines which interfaces are missing from the **show vlans** command output.

For VLAN tagging and Layer 3 subinterfaces, bind an 802.1Q VLAN tag ID to a logical interface.

NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Options For VLANs:

number—A valid VLAN identifier. If you configure multiple VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.

NOTE: If you specify a VLAN identifier, you cannot also use the **all** option. They are mutually exclusive.

all—Specify that the VLAN spans all the VLAN identifiers configured on the member logical interfaces.

NOTE: You cannot specify the **all** option if you include a routing interface in the VLAN.

none—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.

NOTE: Multichassis link aggregation (MC-LAG) does not support the **none** option with the **vlan-id** statement with VLANs.

For 802.1Q Tagging:

number —VLAN tag identifier

Range:

- 1 through 4094 (all switches except EX8200 Virtual Chassis)
- 1 through 4092 (EX8200 Virtual Chassis only)

Default: 1

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

Related Documentation

- *Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support*
- *Example: Configuring a Private VLAN on a Single Switch with ELS Support*
- *Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)*
- *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*
- *Example: Configuring VLANs on Security Devices*
- *Example: Configuring Interfaces and Routing Instances for a User Logical System*
- Rewriting the VLAN Tag on Tagged Frames on page 572
- Binding VLAN IDs to Logical Interfaces on page 251
- **vlan-tagging on page 1445**
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure) for EX Series Switches with ELS support*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

vlan-id-list (Ethernet VLAN Circuit)

Syntax vlan-id-list [*vlan-id vlan-id–vlan-id*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description Binds a single-tag logical interface to a list of VLAN IDs. Configures a logical interface to receive and forward any tag frame whose VLAN ID tag matches the list of VLAN IDs you specify.

NOTE:
When you create a circuit cross-connect (CCC) using VLAN-bundled single-tag logical interfaces on Layer 2 VPN routing instances, the circuit automatically uses **ethernet** encapsulation. For Layer 2 VPN, you need to include the **encapsulation-type** statement and specify the value **ethernet** at either of the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols l2vpn]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]**

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *Junos OS VPNs Library for Routing Devices*.

Options **[*vlan-id vlan-id–vlan-id*]**—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring **vlan-id-list** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
```



```
unit number {
    vlan-id-range 1-4094;
}

[edit interfaces interface-name]
unit 0;
```

Required Privilege Level

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

Related Documentation

- Binding VLAN IDs to Logical Interfaces on page 251
- **encapsulation (Logical Interface) on page 1131**
- **encapsulation on page 1135**
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*
- **flexible-vlan-tagging on page 1169**
- **vlan-tagging on page 1445**
- **vlan-tags (Dual-Tagged Logical Interface) on page 1448**

vlan-id-list (Interface in Bridge Domain)

Syntax vlan-id-list [*number number-number*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 15.1.

Description Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.

Options ***number number***—Individual VLAN IDs separated by a space.

number-number—Starting VLAN ID and ending VLAN ID in an inclusive range.

Range: 1 through 4095

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

Related Documentation

- Configuring a Logical Interface for Trunk Mode on page 262
- Configuring the VLAN ID List for a Trunk Interface on page 262
- *Tunnel Services Overview*
- *Tunnel Interface Configuration on MX Series Routers Overview*

vlan-id-range

Syntax vlan-id-range *vlan-id–vlan-id*

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Bind a range of VLAN IDs to a logical interface.

Options **number**—The first number is the lowest VLAN ID in the range the second number is the highest VLAN ID in the range.

Range: 1 through 4094

NOTE: Configuring **vlan-id-range** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id-range 1-4094;
}

[edit interfaces interface-name]
unit 0;
```

VLAN ID 0 is reserved for tagging the priority of frames.

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

Related Documentation

- *Binding a Range of VLAN IDs to a Logical Interface*

vlan-rewrite

Syntax vlan-rewrite translate (200 500 | 201 501)

Hierarchy Level [edit interfaces *interface-name* unit *number* family bridge interface-mode trunk]
[edit interfaces *interface-name* unit *number* family ethernet-switching interface-mode trunk]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.

Options **translate 200 500**—Translates incoming packets with VLAN 200 to 500.

translate 201 501—Translates incoming packets with VLAN 201 to 501.

translate 202 502—Translates incoming packets with VLAN 202 to 502.

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

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag on page 577

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)

Syntax
vlan-rule (high-low | odd-even);

Hierarchy Level
[edit chassis fpc *slot* pic *slot* forwarding-mode vlan-steering]

Release Information
Statement introduced in Junos OS Release 10.4.

Description
Configure the interoperation mode of the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) when interoperating with 100 gigabit Ethernet interfaces from other vendors.

If no VLAN rule is configured, all tagged packets are distributed to PFE0.

Options
high-low—VLAN IDs 1 through 2047 are distributed to PFE0 and VLAN IDs 2048 through 4096 are distributed to PFE1.

odd-even—Odd number VLAN IDs are distributed to PFE1 and even number VLAN IDs are distributed to PFE0.

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

Related Documentation

- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 448
- **forwarding-mode (100-Gigabit Ethernet) on page 1174**
- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1444**

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)

Syntax

```
vlan-steering {
  vlan-rule (high-low | odd-even);
}
```

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with 100 gigabit Ethernet interfaces from other vendors.

The other statement is 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 VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 448
- **forwarding-mode (100-Gigabit Ethernet) on page 1174**
- **sa-multicast (100-Gigabit Ethernet) on page 1350**
- **vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 1443**

vlan-tagging

Syntax vlan-tagging;

Syntax (QFX Series, NFX Series, and EX4600) vlan-tagging;

Syntax (SRX Series Interfaces) vlan-tagging native-vlan-id *vlan-id*;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

QFX Series, NFX Series, and EX4600 Interfaces [edit interfaces (QFX Series) *interface-name*]
[edit interfaces (QFX Series) interface-range *interface-range-name*]

SRX Series Interfaces [edit interfaces *interface*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

NOTE: For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.

On EX Series switches except for EX4300 and EX9200 switches, the **vlan-tagging** and **family ethernet-switching** statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default **family** setting.

Default VLAN tagging is disabled by default.

SRX Series [Warning: element unresolved in stylesheets: <title> (in <config-options>). This is probably a new element that is not yet supported in the stylesheets.]

SRX Series

native-vlan-id—Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

NOTE: The **native-vlan-id** can be configured only when either **flexible-vlan-tagging** mode or **interface-mode** trunk is configured.

Required Privilege Level

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

Related Documentation

- 802.1Q VLANs Overview on page 244
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- Configuring Tagged Aggregated Ethernet Interfaces on page 138
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *vlan-id*
- *Configuring a Layer 3 Logical Interface*
- *Configuring VLAN Tagging*

vlan-tags

Syntax vlan-tags outer [*tpid*].*vlan-id* [inner [*tpid*].*vlan-id*];

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.
VLAN demux interface support introduced in Junos OS Release 10.2.

Description For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level.

NOTE: The **inner-range** ***vid1–vid2*** option is supported on IQE PICs only.

Options **inner [*tpid*].*vlan-id***—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-vlan-id** predefined variable to dynamically obtain the VLAN ID.

NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range** ***tpid*. *vid1—vid2*** option with the **vlan-tags** statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer [*tpid*].*vlan-id*—A TPID (optional) and a valid VLAN identifier in the format *tpid.vlan-id*. When used in the **dynamic-profiles** hierarchy, specify the **$junos-stacked-vlan-id** predefined variable.

Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.

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

Related Documentation
- Configuring Dual VLAN Tags on page 565

vlan-tags (Dual-Tagged Logical Interface)

Syntax vlan-tags inner-list [*vlan-id vlan-id–vlan-id*] outer <*tpid*.>*vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description (MX Series routers only) Binds a dual-tag logical interface to a list of VLAN IDs. Configures the logical interface to receive and forward any dual-tag frame whose inner VLAN ID tag matches the list of VLAN IDs you specify.

NOTE:
To create a circuit cross-connect (CCC) using VLAN-bundled dual-tag logical interfaces on Layer 2 VPN routing instances, you must include the **encapsulation-type** statement and specify the value **ethernet-vlan** at the one of the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols l2vpn]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]**

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *Junos OS VPNs Library for Routing Devices*.

Options **inner-list [*vlan-id vlan-id vlan-id-vlan-id*]**—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer <*tpid*.>*vlan-id*—An optional Tag Protocol ID (TPID) and a valid VLAN ID.

Range: For TPID, specify a hexadecimal value in the format **0x*nnnn***.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring **inner-list** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through

4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-tags outer vid inner-list 1-4094;
}

[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Required Privilege Level

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

Related Documentation

- Binding VLAN IDs to Logical Interfaces on page 251
- encapsulation (Logical Interface) on page 1131
- encapsulation on page 1135
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*.
- flexible-vlan-tagging on page 1169
- vlan-id-list (Ethernet VLAN Circuit) on page 1438
- vlan-tagging on page 1445

vlan-tags (Stacked VLAN Tags)

Syntax vlan-tags inner *tpid.vlan-id* inner-list *value* inner-range *vid1—vid2* outer *tpid.vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description Bind TPIDs and 802.1Q VLAN tag IDs to a logical interface. TPID fields are used to identify the frame as an IEEE 802.1Q-tagged frame.

Options **inner *tpid.vlan-id***—A TPID and a valid VLAN identifier. TPID is a 16-bit field set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame.

Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

inner-list *value*— List or a set of VLAN identifiers.

NOTE: This is supported on MX Series routers with Trio-based FPCs.

inner-range *tpid. vid1–vid2*—Specify a TPID and a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

NOTE: On the network-to-network (NNI) or egress interfaces of provider edge (PE) routers, you cannot configure the **inner-range *tpid. vid1—vid2*** option with the **vlan-tags** statement for ISP-facing interfaces.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames. For PTX Series, VLAN ID 0 is not supported.

NOTE: Configuring **inner-range** with the entire vlan-id range consumes system resources and is not a best practice. The **inner-range** must be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it has the same result as not specifying a range; however, it consumes Packet Forwarding Engine resources such as VLAN lookup table entries, and so on.

The following examples illustrate this further:

```
[edit interfaces  interface-name]
stacked-vlan-tagging;
unit number {
    vlan-tags outer vid inner-range 1-4094;
}

[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

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

Related Documentation

- Configuring Dual VLAN Tags on page 565
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 249
- **stacked-vlan-tagging on page 1372**

wait-to-block-interval

Syntax wait-to-block-interval *number*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 14.2.

Description Enable the Wait to Block (WTB) timer interval when clearing force switch and manual switch commands.

Options *number*—Wait-to-block interval, in seconds.

Range: 5 through 10 s

Default: 5 s

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

west-interface

Syntax

```
west-interface {
  node-id mac-address;
  control-channel channel-name {
    vlan number;
    interface name interface-name
  }
  interface-none
  ring-protection-link-end;
  virtual-control-channel {
    west-interface name;
    east-interface name;
  }
}
```

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

Description Define one of the two interface ports for Ethernet ring protection, the other being defined by the **east-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

NOTE: Always configure this port second, after configuring the **east-interface** statement.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation

- Ethernet Ring Protection Switching Overview on page 221
- Ethernet Ring Protection Using Ring Instances for Load Balancing on page 857
- **east-interface on page 1129**
- **ethernet-ring on page 1153**
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS*
- *Configuring Ethernet Ring Protection Switching on Switches (CLI Procedure)*

CHAPTER 40

Operational Commands

- show ethernet-switching redundancy-groups
- show interfaces (Adaptive Services)
- show interfaces (Aggregated Ethernet)
- show interfaces demux0 (Demux Interfaces)
- show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)
- show interfaces (far-end-interval)
- show interfaces (Fast Ethernet)
- show interfaces
- show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)
- show interfaces (PPPoE)
- show interfaces interface-set (Ethernet Interface Set)
- show interfaces interface-set queue
- show interfaces interval
- show interfaces irb
- show interfaces mac-database
- show interfaces mc-ae
- show interfaces transport pm
- show l2-learning instance
- show l2-learning redundancy-groups
- show lacp interfaces
- show lldp
- show lldp local-information
- show lldp neighbors
- show lldp remote-global-statistics
- show lldp statistics
- show oam ethernet connectivity-fault-management delay-statistics
- show oam ethernet connectivity-fault-management forwarding-state
- show oam ethernet connectivity-fault-management interfaces
- show oam ethernet connectivity-fault-management linktrace path-database
- show oam ethernet connectivity-fault-management loss-statistics
- show oam ethernet connectivity-fault-management mep-database
- show oam ethernet connectivity-fault-management mep-statistics
- show oam ethernet connectivity-fault-management path-database
- show oam ethernet connectivity-fault-management policer
- show oam ethernet connectivity-fault-management sla-iterator-statistics
- show oam ethernet connectivity-fault-management synthetic-loss-statistics

- show oam ethernet evc
- show oam ethernet fnp interface
- show oam ethernet fnp messages
- show oam ethernet fnp status
- show oam ethernet link-fault-management
- show oam ethernet lmi
- show oam ethernet lmi statistics
- show pppoe interfaces
- show pppoe service-name-tables
- show pppoe sessions
- show pppoe statistics
- show pppoe underlying-interfaces
- show pppoe version
- show protection-group ethernet-ring aps
- show protection-group ethernet-ring configuration
- show protection-group ethernet-ring data-channel
- show protection-group ethernet-ring flush-info
- show protection-group ethernet-ring interface
- show protection-group ethernet-ring node-state
- show protection-group ethernet-ring statistics
- show protection-group ethernet-ring vlan
- show security macsec connections (MX Series)
- show security macsec statistics (MX Series)
- show security mka sessions (MX Series)
- show security mka statistics (MX Series)
- show vrrp
- traceroute ethernet

clear interfaces interface-set statistics

Syntax clear interfaces interface-set statistics *interface-set-name*

Release Information Command introduced in Junos OS Release 8.5.

Description Set interface set statistics to zero.

Options ***interface-set-name***—Set statistics on a specified interface set to zero. Wildcard values can be used in the interface set name. This command will not clear the statistics of the member logical interfaces.

Required Privilege Level clear

List of Sample Output clear interfaces interface-set statistics on page 1456

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear interfaces interface-set statistics

`user@host>` **clear interfaces interface-set statistics**

clear interfaces interval

Syntax clear interfaces interval *interface-name*

Release Information Command introduced before Junos OS Release 7.4.

Description Clear the channel service unit (CSU) alarm and defect counters so that only the current time interval is displayed. This operation affects the **show interface interval** command, but not an SNMP query.

Options *interface-name*—Name of a particular interface.

Required Privilege Level clear

Related Documentation

- show interfaces interval on page 1974

List of Sample Output clear interfaces interval on page 1457

Output Fields See show interfaces interval for an explanation of output fields.

Sample Output

clear interfaces interval

The following example displays the output for a T3 interface before and after the **clear interfaces** command is entered:

```
user@host> show interfaces interval t3-0/3/0:4
Physical interface: t3-0/3/0:4, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  16:28-16:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
CES: 195, CSES: 195, SEFS: 195, UAS: 206
  14:58-15:13:
    LCV: 35, PCV: 163394, CCV: 54485, LES: 0, PES: 35, PSES: 35, CES:
35, CSES: 35, SEFS: 35, UAS: 32
```



```
  Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

`user@host>` **clear interfaces interval t3-0/3/0:4**

```
user@host> show interfaces interval t3-0/3/0:4
Physical interface: t3-0/3/0:4, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
  Interval Total:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0, SEFS: 0,
UAS: 0
```

clear interfaces aeX forwarding-options load-balance state

Syntax clear interfaces ae*X* unit *logical-unit-number* aggregate forwarding-options load-balance state

Release Information Command introduced in Junos OS Release 13.2R1.

Description Clear the specified aggregate Ethernet interface load balancing state and re-create it newly. If the traffic flows become aged frequently, then the device needs to remove or refresh the load balancing states. As a result, you must configure rebalancing or run the clear command at periodic intervals for proper load-balancing. Otherwise, traffic skewing can occur.

If you observe load distribution to be not very effective, you can clear the load-balancing states or use rebalancing functionality to cause an automatic clearance of the hardware states. When you configure the rebalancing facility, traffic flows can get redirected to different links, which can cause packet reordering.

Options **ae*X***—Name of a particular aggregated Ethernet interface.

logical-unit-number—Number of the logical unit of the interface.

forwarding-options load-balance state—Cause the load-balancing state to be cleared for the specific interface.

Required Privilege Level clear

Related Documentation

- show interfaces interval on page 1974

List of Sample Output clear interfaces aeX aggregate forwarding-options on page 1460

Sample Output

clear interfaces aeX aggregate forwarding-options

```
user@host> clear interfaces ae1 aggregate forwarding-options load-balance state
```


clear interfaces aggregate forwarding-options load-balance state

Syntax clear interfaces aggregate forwarding-options load-balance state

Release Information Command introduced in Junos OS Release 13.2R1.

Description Clear all the aggregate Ethernet interface load balancing states and re-create them newly. If the traffic flows become aged frequently, then the device needs to remove or refresh the load balancing states. As a result, you must configure rebalancing or run the clear command at periodic intervals for proper load-balancing. Otherwise, traffic skewing can occur.

Options ***interface-name***—Name of a particular interface.

Required Privilege Level clear

Related Documentation

- show interfaces interval on page 1974

List of Sample Output clear interfaces aggregate forwarding-options on page 1461

Sample Output

clear interfaces aggregate forwarding-options

```
user@host> clear interfaces aggregate forwarding-options load-balance state
```

clear interfaces transport pm

Syntax

clear interfaces transport pm (all | optics | otn) (all | current | current-day) (all | *interface-name*)

Release Information

Command introduced in Junos OS Release 14.2 on the PTX Series.

Description

Clear optics and OTN information from the transport performance monitoring data.

Options

(all | optics | otn)—Clear both optics and OTN information or either only optics or only OTN information.

(all | current | current-day)—Clear information for the current 15-minute interval, the ninety-six 15-minute intervals, the current day, and the previous day; information only for the current 15-minute interval; or information only for the current 24 hours.

(all | *interface-name*)—Clear information for all interfaces or only for the specified interface (for example, *et-fpc/pic/port*)).

Required Privilege Level

clear

Related Documentation

- show interfaces transport pm on page 1994
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 475
- tca on page 974

List of Sample Output

clear interfaces transport pm on page 1462

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear interfaces transport pm

```
user@host> clear interfaces transport pm transport otn current all
```

clear lldp neighbors

Syntax

clear lldp neighbor
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Clear information regarding all Link Layer Discovery Protocol (LLDP) neighbors or LLDP neighbors of the specified interface.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about interface names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

interface *interface-name*—(Optional) Clear the LLDP neighbors on the specified interface.

Required Privilege Level

clear

Related Documentation

- clear lldp statistics on page 1464

List of Sample Output

clear lldp neighbors on page 1463
clear lldp neighbors interface ge-0/1/1.0 on page 1463

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show lldp neighbors** command before and after clearing the LLDP neighbors to verify the clear operation.

Sample Output

clear lldp neighbors

```
user@switch> clear lldp neighbors
```

clear lldp neighbors interface ge-0/1/1.0

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```


clear lldp statistics

Syntax

clear lldpp neighbor
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Clear all Link Layer Discovery Protocols (LLDP) statistics or LLDP statistics associated with the specified interface.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about interface names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

interface *interface-name*—(Optional) Clear LLDP statistics on the specified interface.

Required Privilege Level

clear

Related Documentation

- clear lldp neighbors on page 1463

List of Sample Output

clear lldp statistics on page 1464
clear lldp statistics interface ge-0/1/1.0 on page 1464

Output Fields

When you enter this command, you are provided no feedback on the status of your request. You can enter the **show lldp statistics** command before and after clearing the LLDP statistics to verify the clear operation.

Sample Output

clear lldp statistics

```
user@switch> clear lldp statistics
```

clear lldp statistics interface ge-0/1/1.0

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

clear oam ethernet connectivity-fault-management continuity-measurement

Syntax

clear oam ethernet connectivity-fault-management continuity-measurement
maintenance-domain ***md-name***
maintenance-association ***ma-name***
<**local-mep** ***local-mep-id***>
<**remote-mep** ***remote-mep-id***>

Release Information

Command introduced in Junos OS Release 11.1.

Description

For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear the existing continuity measurement and restart counting the operational uptime (that is, the total time during which CCM adjacency is active for a particular remote MEP.).

Options

maintenance-domain ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

local-mep ***local-mep-id***—(Optional) Display connectivity fault management information for the specified local MEP only.

remote-mep ***remote-mep-id***—(Optional) Display connectivity fault management information for the specified remote MEP only.

Required Privilege Level

view

Related Documentation

- Managing Continuity Measurement Statistics on page 803
- *Ethernet Interfaces Feature Guide for Routing Devices*

List of Sample Output

clear oam ethernet connectivity-fault-management continuity-measurement on page 1465

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management continuity-measurement

```
user@host> clear oam ethernet connectivity-fault-management continuity-measurement
maintenance-domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
Continuity measurement restarted.
```

clear oam ethernet connectivity-fault-management linktrace path-database

Syntax **clear oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-association** ***ma-name*** **maintenance-domain** ***md-name***

Release Information Command introduced in Junos OS Release 9.0.

Description Clear all the linktrace entries and the relevant path information from the database for a particular remote host on M320, MX Series, T320, and T640 routers.

Options

mac-address—Clear connectivity fault management path database information for the specified MAC address of the remote host.

maintenance-association ***ma-name***—Clear connectivity fault management path database information for the specified maintenance association.

maintenance-domain ***md-name***—Clear connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level view

Sample Output

clear oam ethernet connectivity-fault-management linktrace path-database

```
user@host> clear oam ethernet connectivity-fault-management linktrace path-database
maintenance-domain md1 maintenance-association ma3 00058573e483
This command produces no output.
```


clear oam ethernet connectivity-fault-management loss-statistics

Syntax

clear oam ethernet connectivity-fault-management loss-statistics
<interface *ethernet-interface-name*>
<level *md-level*>

Release Information

Command introduced in Junos OS Release 11.1.

Description

For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear all loss statistics maintained by CFM for a given maintenance domain and maintenance association.

In addition, for Ethernet interfaces on MX Series routers , clear any ITU-T Y.1731 Ethernet frame loss measurement (ETH-LM) statistics.

By default, the command clears ETH-LM statistics for CFM maintenance association end points (MEPs) attached to any interface on the router.

Options

interface *ethernet-interface-name*—(Optional) Clear ETH-LM statistics and ETH-LM frame counts only for MEPs attached to the specified Ethernet physical interface.

level *md-level*—(Optional) Clear ETH-LM statistics and ETH-LM frame counts only for MEPs within CFM maintenance domains (MDs) of the specified level.

Required Privilege Level

view

Related Documentation

- Clearing ETH-LM Statistics on page 803
- Displaying ETH-LM Statistics on page 802
- Managing ETH-LM Statistics on page 802

List of Sample Output

clear oam ethernet connectivity-fault-management loss-statistics on page 1467

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management loss-statistics

```
user@host> clear oam ethernet connectivity-fault-management loss-statistics
Cleared loss measurements statistics of all CFM sessions
```


clear oam ethernet connectivity-fault-management policer

Syntax clear oam ethernet connectivity-fault-management policer maintenance-domain *md-name* maintenance-association *ma-name*

Release Information Command introduced in Junos OS Release 10.0.

Description On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, M120, MX Series, T320, and T640 routers, clear connectivity-fault-management policer statistics.

Options The following options are supported:

maintenance-domain md-name—Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are cleared for all maintenance associations for all maintenance domains.

maintenance-association ma-name—Name of an existing CFM maintenance association. If this option is not specified, policer statistics are cleared for all maintenance associations for given maintenance domain. This option cannot be specified without specifying maintenance-domain name.

Required Privilege Level view

Related Documentation

- show oam ethernet connectivity-fault-management policer on page 2080

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear oam ethernet connectivity-fault-management policer

```
user@host> clear oam ethernet connectivity-fault-management policer
Policer statistics cleared
```

clear oam ethernet connectivity-fault-management policer maintenance-domain md-name maintenance-association ma-name

```
user@host> clear oam ethernet connectivity-fault-management policer
maintenance-domain md5 maintenance-association ma5-1
Policer statistics cleared
```


clear oam ethernet connectivity-fault-management statistics

Syntax

clear oam ethernet connectivity-fault-management statistics
<interface *ethernet-interface-name*>
<level *md-level*>

Release Information

Command introduced in Junos OS Release 8.4.
Support for ETH-DM statistics and frame counts added in Junos OS Release 9.5.

Description

For all routers that support IEEE 802.1ag OAM connectivity-fault management (CFM), clear all statistics maintained by CFM.

In addition, for Ethernet interfaces on Dense Port Concentrators (DPCs) in MX Series routers only, also clear any ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) statistics and ETH-DM frame counts.

By default, the command clears CFM statistics and ETH-DM statistics and frame counts for CFM maintenance association end points (MEPs) attached to any interface on the router.

Options

ethernet-interface-name—(Optional) Clear CFM statistics, ETH-DM statistics, and ETH-DM frame counts only for MEPs attached to the specified Ethernet physical interface.

level—(Optional) Clear CFM statistics, ETH-DM statistics, and ETH-DM frame counts only for MEPs within CFM maintenance domains (MDs) of the specified level.

Required Privilege Level

view

Related Documentation

- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- show oam ethernet connectivity-fault-management interfaces on page 2037
- show oam ethernet connectivity-fault-management mep-database on page 2055
- show oam ethernet connectivity-fault-management mep-statistics on page 2066

List of Sample Output

clear oam ethernet connectivity-fault-management statistics on page 1469

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management statistics

```
user@host> clear oam ethernet connectivity-fault-management statistics
Cleared statistics of all CFM sessions
```

clear oam ethernet connectivity-fault-management synthetic-loss-measurement

Syntax

clear oam ethernet connectivity-fault-management synthetic-loss-measurement
maintenance-domain ***md-name***
maintenance-association ***ma-name***
<**local-mep** ***local-mep-id***>
<**remote-mep** ***remote-mep-id***>

Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.

Description

(MX Series routers)—For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear the existing on-demand Ethernet synthetic loss measurement (ETH-SLM) statistics and restart counting the ETH-SLM frame counts and statistics.

Options

maintenance-domain ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

local-mep ***local-mep-id***—(Optional) Clear connectivity fault management information for the specified local MEP only.

remote-mep ***remote-mep-id***—(Optional) Clear connectivity fault management information for the specified remote MEP only.

Required Privilege Level

view

Related Documentation

- monitor ethernet synthetic-loss-measurement on page 1487
- show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 2088

List of Sample Output

clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 1471

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management synthetic-loss-measurement

```
user@host> clear oam ethernet connectivity-fault-management synthetic-loss-measurement
maintenance-domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
Synthetic loss measurement restarted.
```

clear oam ethernet link-fault-management state

Syntax clear oam ethernet link-fault-management state <*interface-name*>

Release Information Command introduced in Junos OS Release 8.4.

Description On all M Series, MX Series, ACX series, PTX Series, T320, and T640 routers, clear link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on Ethernet interfaces.

Options **none**—Clear OAM link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on all Ethernet interfaces.

interface-name—(Optional) Clear OAM link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on the specified Ethernet interface only.

Required Privilege Level view

List of Sample Output clear oam ethernet link-fault-management state on page 1472

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet link-fault-management state

```
user@host> clear oam ethernet link-fault-management state ge-0/3/3
Cleared link-fault-management state for interface ge-0/3/3
```


clear oam ethernet link-fault-management statistics

Syntax clear oam ethernet link-fault-management <*interface-name*>

Release Information Command introduced in Junos OS Release 8.2.

Description On M320, M120, MX Series, PTX Series, T320, and T640 routers, clear Operation, Administration, and Management (OAM) link fault management statistics or state information from Ethernet interfaces.

Options **none**—Clear OAM link fault management statistics from all Ethernet interfaces.

interface-name—(Optional) Clear OAM link fault management statistics from the specified Ethernet interface only.

Required Privilege Level view

List of Sample Output clear oam ethernet link-fault-management statistics on page 1473

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet link-fault-management statistics

```
user@host> clear oam ethernet link-fault-management statistics
Cleared link-fault-management statistics for all interfaces
```


clear protection-group ethernet-ring statistics

Syntax

clear protection-group ethernet-ring statistics
<group-name *group-name*>

Release Information

Command introduced in Junos OS Release 9.4.

Description

On MX Series routers, clear the statistics for all Ethernet ring protection groups or a specific Ethernet ring protection group.

Options

group-name ***group-name***—(Optional) Clear the Ethernet ring protection statistics for the specified group.

Required Privilege Level

view

List of Sample Output

clear protection-group ethernet-ring statistics on page 1474
clear protection-group ethernet-ring statistics on page 1474

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear protection-group ethernet-ring statistics

To clear all Ethernet ring protection group statistics for all protection groups, use the following command:

`user@host>` **clear protection-group ethernet-ring statistics**

Sample Output

clear protection-group ethernet-ring statistics

To clear Ethernet ring protection group statistics for the group *my_prot_group*, use the following command:

`user@host>` **clear protection-group ethernet-ring statistics group-name *my_prot_group***

clear security mka statistics (MX Series)

Syntax

clear security mka statistics
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description

Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.

You are clearing the statistics that are viewed using the **show security mka statistics** when you enter this command.

Options

none—Clear all MKA counters for all interfaces on the switch.

interface ***interface-name***—(Optional) Clear MKA traffic counters for the specified interface only.

Required Privilege Level

clear

Sample Output

clear security mka statistics

```
user@switch> clear security mka statistics
```

clear security mka statistics (MX Series)

Syntax

clear security mka statistics
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.

Description

Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.

You are clearing the statistics that are viewed using the **show security mka statistics** when you enter this command.

Options

none—Clear all MKA counters for all interfaces on the switch.

interface ***interface-name***—(Optional) Clear MKA traffic counters for the specified interface only.

Required Privilege Level

clear

Sample Output

clear security mka statistics

```
user@switch> clear security mka statistics
```

monitor ethernet delay-measurement

Syntax

monitor ethernet delay-measurement
maintenance-domain *md-name*
maintenance-association *ma-name*
(one-way | two-way)
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<local-mep *mep-id*>
<wait *interval-seconds*>
<priority *802.1p value*>
<size>
<no-session-id-tlv>
<xml>

Release Information

Command introduced in Junos OS Release 9.5.
local-mep option introduced in Junos OS Release 15.1.

Description

Start an ITU-T Y.1731 Ethernet frame delay measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. Frame delay measurement statistics are stored at one of the MEPs for later retrieval.

NOTE: **If you attempt to monitor delays to a nonexistent MAC address, you must type Ctrl +c to explicitly quit the monitor ethernet delay-measurement command and return to the CLI command prompt.**

To start an Ethernet frame delay measurement session, the router initiates an exchange of frames carrying one-way or two-way frame delay measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet frame delay measurement PDU frames exchanged to measure frame delay times—are displayed as the runtime output of the **monitor ethernet delay-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet frame delay measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay—The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

Options **maintenance-domain** ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

one-way—Measurement type is one-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a one-way delay measurement request (1DM) frame and the time at which the receiver MEP receives the frame.

two-way—Measurement type is two-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a two-way delay measurement message (DMM) frame and the time at which the initiator MEP receives an associated two-way delay measurement reply (DMR) frame from the responder MEP, subtracting the time elapsed at the responder MEP.

mep ***remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet frame delay measurement. The discovered MAC address of the peer MEP is used. The range of values is 1 through 8191.

remote-mac-address—Unicast MAC address of the peer MEP with which to perform Ethernet frame delay measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**. Multicast MAC addresses are not supported.

count ***frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is 1 through 65,535 frames. The default value is 10 frames.

local-mep ***mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance association end point.

wait ***interval-seconds***—(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

priority ***802.1p value***—(Optional) Priority of the delay measurement request frame supported by both one-way delay measurement and two-way delay measurement. The range of values is from 0 through 7. The default value is zero.

size—(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

no-session-id-tlv—(Optional) Prevent insertion of the session ID TLV in the request frame.

xml—(Optional) Allow the output of the command to be displayed in XML format supported by both one-way delay measurement and two-way delay measurement. Note that the only way to get output in XML format is to use the **xml** argument. The **display xml** command does not work.

Additional Information

To display the frame counts collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer:

- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer.

- **show oam ethernet connectivity-fault-management delay-statistics**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management statistics** command, described in the CLI Explorer.

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

Required Privilege Level

trace and maintenance

List of Sample Output

monitor ethernet delay-measurement one-way on page 1480
monitor ethernet delay-measurement two-way on page 1481
monitor ethernet delay-measurement two-way (Invalid DMR Frames Received) on page 1481

Output Fields

The **monitor ethernet delay-measurement** command displays different output at the CLI, depending on whether you start a one-way or two-way frame delay measurement:

- Table 102 on page 1479 lists the run-time output fields for the **monitor ethernet delay-measurement one-way** command.
- Table 103 on page 1480 lists the run-time output fields for the **monitor ethernet delay-measurement two-way** command.

Output fields are listed in the approximate order in which they appear.

Table 102: monitor ethernet delay-measurement one-way Output Fields

Output Field Name	Output Field Description
One-way ETH-DM request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
1DM Frames sent	PDU frames sent to the remote MEP in this ETH-DM session.
Packets transmitted	Total number of 1DM PDU frames sent to the remote MEP during this measurement session.

Table 102: monitor ethernet delay-measurement one-way Output Fields (continued)

Output Field Name	Output Field Description
Average delay	Average two-way frame delay measured in this session.
Average delay variation	Average frame jitter measured in this session.
Best case delay	Lowest two-way frame delay measured in this session.
Worst case delay	Highest two-way frame delay measured in this session.
NOTE: For one-way delay measurement, these CLI output fields display **NA** ("not applicable") at the initiator MEP because one-way frame delay measurements occur at the receiver MEP.	

Table 103: monitor ethernet delay-measurement two-way Output Fields

Output Field Name	Output Field Description
Two-way Ethernet frame delay measurement request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
DMR received from	Unicast MAC address of the remote MEP that transmitted this DMR frame in response to a DMM frame.
Delay	Two-way delay, in microseconds, for the initiator-transmitted DMM frame.
Delay variation	Difference, in microseconds, between the current and previous delay values. This is also known as frame jitter.
Packets transmitted	Total number of DMM PDU frames sent to the remote MEP in this measurement session.
Valid packets received	Total number of DMR PDU frames received from the remote MEP in this measurement session.
Average delay	Average two-way frame delay measured in this session.
Average delay variation	Average frame jitter measured in this session.
Best case delay	Lowest two-way frame delay measured in this session.
Worst case delay	Highest two-way frame delay measured in this session.

Sample Output

monitor ethernet delay-measurement one-way

`user@host>` **monitor ethernet delay-measurement one-way 00:05:85:73:39:4a maintenance-domain md6 maintenance-association ma6 count 10**


```
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

monitor ethernet delay-measurement two-way

user@host> **monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain md6 maintenance-association ma6 count 10**

```
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a  Delay: 100 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 111 usec  Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a  Delay: 110 usec  Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a  Delay: 119 usec  Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a  Delay: 122 usec  Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 108 usec  Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

monitor ethernet delay-measurement two-way (Invalid DMR Frames Received)

user@host> **monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain md6 maintenance-association ma6 count 10**

```
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a  Delay: 100 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a  Delay: 111 usec  Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a  Delay: 110 usec  Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a  Delay: 119 usec  Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a  Delay: 122 usec  Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a  Delay: 92 usec   Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a with invalid timestamp(s).
DMR received from 00:05:85:73:39:4a  Delay: 108 usec  Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 9, Invalid packets received: 1
Average delay: 105 usec, Average delay variation: 9 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

monitor ethernet loss-measurement

Syntax

monitor ethernet loss-measurement
maintenance-domain *md-name*
maintenance-association *ma-name*
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<local-mep *mep-id*>
<wait *interval-seconds*>
<priority *802.1p value*>
<no-session-id-tlv>
<xml>

Release Information

Command introduced in Junos OS Release 11.1.
local-mep option introduced in Junos OS Release 15.1

Description

Start an ITU-T Y.1731 Ethernet frame loss measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a count of transmitted and received data frames between the pair of MEPs. Frame loss measurement statistics are stored at one of the MEPs for later retrieval. For MX Series routers, supports point-to-point down MEPs for Ethernet interfaces (as per IEEE 802.1ag over VPWS).

NOTE: If you attempt to monitor loss to a nonexistent MAC address, you must type Ctrl + c to explicitly quit the **monitor ethernet loss-measurement** command and return to the CLI command prompt.

To start an Ethernet frame loss measurement session, the router first sends frames with ETH-LM information to a peer MEP and similarly receives frames with ETH-LM information from the peer MEP. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the **monitor ethernet loss-measurement** command and are also stored at the initiator. The frames counts are stored at both the initiator and the receiver MEPs for later retrieval.

Options

maintenance-domain ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

mep ***remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8192.

remote-mac-address—Unicast MAC address of the peer MEP with which to perform Ethernet frame loss measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn*** (for example, **0011.2233.4455** or **00:11:22:33:44:55**). Multicast MAC addresses are not supported.

count ***frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65535 frames. The default value is 10 frames.

local-mep ***mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

wait ***interval-seconds***—(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

priority ***802.1p value***—(Optional) Priority of the delay measurement request frame. The range of values is from 0 through 7. The default value is 1 second.

no-session-id-tlv—(Optional) Disable the **session id TLV** argument set in the request frame.

xml—(Optional) Allow the output of the command to be displayed in XML format.

Additional Information

To display the interator output for an LM session, run the following command:

- **show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator <profile> maintenance-association <MA> maintenance-domain <MD> local-mep <MEP> remote-mep <RMEP>**

To display the frame counts collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer:

- **show oam ethernet connectivity-fault-management loss-statistics**
- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at an MEP as the result of this command, see the following command descriptions in the CLI Explorer:

- **show oam ethernet connectivity-fault-management delay-statistics**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management loss-statistics maintenance-domain *md-name* maintenance-association *ma-name*** command, as described in the CLI Explorer.

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

Required Privilege Level

trace and maintenance

Related Documentation

- Ethernet Frame Loss Measurement Overview on page 729
- *Junos OS Network Interfaces Library for Routing Devices*
- CLI Explorer

List of Sample Output

monitor ethernet loss-measurement on page 1485

Output Fields

Table 104 on page 1484 lists the output fields for the **monitor ethernet loss-measurement** command and their descriptions. Output fields are listed in the approximate order in which they appear.

Table 104: monitor ethernet loss-measurement output fields

Output Field Name	Output Field Description
Ethernet loss delay measurement request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
LMR received from	Unicast MAC address of the remote MEP that transmitted this LMR frame in response to a loss measurement message (LMM) frame.
Near-end frame loss	Count of frame loss associated with ingress data frames.
Far-end frame loss	Count of frame loss associated with egress data frames.
Near-end loss ratio	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the ingress interface.
Far-end loss ratio	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the egress interface.
LMM packets transmitted	Total number of LMM PDU frames sent to the remote MEP in this measurement session.
LMR packets received	Total number of LMR PDU frames received from the remote MEP in this measurement session.
Average near-end frame loss	Average frame loss measured in this session associated with ingress data frames.
Average near-end loss ratio	Average frame loss ratio measured in this session associated with ingress data frames.
Average far-end frame loss	Average frame loss measured in this session associated with egress data frames.
Average far-end loss ratio	Average frame loss ratio measured in this session associated with egress data frames.
Near-end best case frame loss	Lowest frame loss measured in this session associated with ingress data frames.
Near-end best case loss ratio	Lowest frame loss ratio measured in this session associated with ingress data frames.

Table 104: monitor ethernet loss-measurement output fields (continued)

Output Field Name	Output Field Description
Near-end worst case frame loss	Highest frame loss measured in this session associated with ingress data frames.
Near-end worst case loss ratio	Highest frame loss ratio measured in this session associated with ingress data frames.
Far-end best case frame loss	Lowest frame loss measured in this session associated with egress data frames.
Far-end best case loss ratio	Lowest frame loss ratio measured in this session associated with egress data frames.
Far-end worst case frame loss	Highest frame loss measured in this session associated with egress data frames.
Far-end worst case loss ratio	Highest frame loss ratio measured in this session associated with egress data frames.

Note that in the preceding table, the term *number of service frames not delivered* is the difference between the number of service frames arriving at the ingress Ethernet flow point and the number of service frames delivered at the egress Ethernet flow point in a point-to-point Ethernet connection.

Sample Output

monitor ethernet loss-measurement

user@host> **monitor ethernet loss-measurement mep 2 64:87:88:6a:da:94 maintenance-domain md maintenance-association ma count 10**

```
ETH-LM request to 64:87:88:6a:da:94, Interface ge-2/3/2.0

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):260
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.43537%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51546%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):257
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.31615%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):261
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.77551%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
```



```
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):260
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.43537%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51020%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):259
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.09524%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):519
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.71795%

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0           Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51020%
Near-end frame loss(EIR)      :0           Far-end frame loss(EIR):259
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.09524%

--- Loss measurement statistics ---
LMM packets transmitted: 10, Valid LMR packets received: 8
Average near-end loss(CIR)         : 0.00000
Average near-end loss ratio(CIR)   : 0.00000%
Average far-end loss(CIR)          : 0.42857
Average far-end loss ratio(CIR)    : 0.21941%
Near-end best case loss(CIR)       : 0
Near-end best case loss ratio(CIR) : 0.00000%
Near-end worst case loss(CIR)      : 0
Near-end worst case loss ratio(CIR): 0.00000%
Far-end best case loss(CIR)        : 0
Far-end best case loss ratio(CIR)  : 0.00000%
Far-end worst case loss(CIR)       : 1
Far-end worst case loss ratio(CIR) : 0.51546%
Average near-end loss(EIR)         : 0.00000
Average near-end loss ratio(EIR)   : 0.00000%
Average far-end loss(EIR)          : 296.42857
Average far-end loss ratio(EIR)    : 88.41011%
Near-end best case loss(EIR)       : 0
Near-end best case loss ratio(EIR) : 0.00000%
Near-end worst case loss(EIR)      : 0
Near-end worst case loss ratio(EIR): 0.00000%
Far-end best case loss(EIR)        : 257
Far-end best case loss ratio(EIR)  : 88.09524%
Far-end worst case loss(EIR)       : 519
Far-end worst case loss ratio(EIR) : 88.77551%
```


monitor ethernet synthetic-loss-measurement

Syntax

monitor ethernet synthetic-loss-measurement
maintenance-domain *md-name*
maintenance-association *ma-name*
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<local-mep *mep-id*>
<wait *interval-milliseconds*>
<priority *802.1p value*>
<size>
<xml>

Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.
local-mep option introduced in Junos OS Release 15.1

Description

(MX Series routers) Start an ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. ETH-SLM (also called synthetic frame loss measurement) statistics are stored at one of the MEPs for later retrieval.

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must press Ctrl +c to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.

To start an Ethernet synthetic frame loss measurement session, the router initiates an exchange of frames carrying synthetic frame loss measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet synthetic frame loss measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time output of the **monitor ethernet synthetic-loss-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet synthetic frame loss measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay—The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

Options

maintenance-domain *md-name*—Name of an existing CFM maintenance domain.

maintenance-association *ma-name*—Name of an existing CFM maintenance association.

mep ***remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet synthetic frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8191.

remote-mac-address—Unicast MAC address of the peer MEP with which to perform Ethernet synthetic frame loss measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn*:*nn*:*nn*:*nn*:*nn*:*nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**. Multicast MAC addresses are not supported.

count ***frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65,535 frames. The default value is 10 frames.

local-mep ***mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

wait ***interval-milliseconds***—(Optional) Number of milliseconds to wait between sending frames. You must specify this value in multiples of 100 milliseconds. The range of values is from 100 through 50000 milliseconds. The default value is 100 milliseconds.

priority ***802.1p value***—(Optional) Priority of the ETH-SLM request frame supported. The range of values is from 0 through 7. The default value is zero.

size—(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

xml—(Optional) Allow the output of the command to be displayed in XML format for ETH-SLM. Note that the only way to get output in XML format is to use the **xml** argument. The **display xml** command does not work.

Additional Information

To display the frame counts collected at a MEP as the result of this command, use the following commands as described in the CLI Explorer:

- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at a MEP as the result of this command, use the following commands as described in the CLI Explorer.

- **show oam ethernet connectivity-fault-management synthetic-loss-measurement**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management statistics** command, described in the CLI Explorer.

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

Required Privilege Level

trace and maintenance

Related Documentation

- clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 1471
- monitor ethernet synthetic-loss-measurement on page 1487
- show oam ethernet connectivity-fault-management loss-statistics on page 2051

List of Sample Output

monitor ethernet synthetic-loss-measurement on page 1490

Output Fields

The **monitor ethernet synthetic-loss-measurement** command displays different output at the CLI, depending on when you start a synthetic frame loss measurement:

- Table 105 on page 1489 lists the run-time output fields for the **monitor ethernet synthetic-loss-measurement** command.

Output fields are listed in the approximate order in which they appear.

Table 105: monitor ethernet synthetic-loss-measurement Output Fields

Output Field Name	Output Field Description
ETH-SLM request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames received by the source MEP from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFCl at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t	Value of the local counter RxFCl at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

monitor ethernet synthetic-loss-measurement

user@host> **monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a maintenance-domain md6 maintenance-association ma6 count 10**

```
ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0
      Synthetic Loss measurement statistics:
       SLM packets sent                              : 100
       SLR packets received                          : 100
       Accumulated SLM statistics:
       Local TXFC1 value                             : 100
       Local RXFC1 value                             : 100
       Last Received SLR frame TXFCf(tc)             : 100
       Last Received SLR frame TXFCb(tc)             : 100
       SLM Frame Loss:
       Frame Loss (far-end)                          : 0 (0.00 %)
       Frame Loss (near-end)                         : 0 (0.00 %)
```

monitor ethernet synthetic-loss-measurement

Syntax

monitor ethernet synthetic-loss-measurement
maintenance-domain *md-name*
maintenance-association *ma-name*
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<local-mep *mep-id*>
<wait *interval-milliseconds*>
<priority *802.1p value*>
<size>
<xml>

Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.
local-mep option introduced in Junos OS Release 15.1

Description

(MX Series routers) Start an ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. ETH-SLM (also called synthetic frame loss measurement) statistics are stored at one of the MEPs for later retrieval.

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must press Ctrl +c to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.

To start an Ethernet synthetic frame loss measurement session, the router initiates an exchange of frames carrying synthetic frame loss measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet synthetic frame loss measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time output of the **monitor ethernet synthetic-loss-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet synthetic frame loss measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay—The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

Options

maintenance-domain *md-name*—Name of an existing CFM maintenance domain.

maintenance-association *ma-name*—Name of an existing CFM maintenance association.

mep ***remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet synthetic frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8191.

remote-mac-address—Unicast MAC address of the peer MEP with which to perform Ethernet synthetic frame loss measurement. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn*:*nn*:*nn*:*nn*:*nn*:*nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**. Multicast MAC addresses are not supported.

count ***frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65,535 frames. The default value is 10 frames.

local-mep ***mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

wait ***interval-milliseconds***—(Optional) Number of milliseconds to wait between sending frames. You must specify this value in multiples of 100 milliseconds. The range of values is from 100 through 50000 milliseconds. The default value is 100 milliseconds.

priority ***802.1p value***—(Optional) Priority of the ETH-SLM request frame supported. The range of values is from 0 through 7. The default value is zero.

size—(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

xml—(Optional) Allow the output of the command to be displayed in XML format for ETH-SLM. Note that the only way to get output in XML format is to use the **xml** argument. The **display xml** command does not work.

Additional Information

To display the frame counts collected at a MEP as the result of this command, use the following commands as described in the CLI Explorer:

- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at a MEP as the result of this command, use the following commands as described in the CLI Explorer.

- **show oam ethernet connectivity-fault-management synthetic-loss-measurement**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management statistics** command, described in the CLI Explorer.

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

Required Privilege Level

trace and maintenance

Related Documentation

- clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 1471
- monitor ethernet synthetic-loss-measurement on page 1487
- show oam ethernet connectivity-fault-management loss-statistics on page 2051

List of Sample Output

monitor ethernet synthetic-loss-measurement on page 1494

Output Fields

The **monitor ethernet synthetic-loss-measurement** command displays different output at the CLI, depending on when you start a synthetic frame loss measurement:

- Table 105 on page 1489 lists the run-time output fields for the **monitor ethernet synthetic-loss-measurement** command.

Output fields are listed in the approximate order in which they appear.

Table 106: monitor ethernet synthetic-loss-measurement Output Fields

Output Field Name	Output Field Description
ETH-SLM request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames received by the source MEP from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFCl at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t	Value of the local counter RxFCl at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

monitor ethernet synthetic-loss-measurement

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0
      Synthetic Loss measurement statistics:
       SLM packets sent                            : 100
       SLR packets received                        : 100
       Accumulated SLM statistics:
       Local TXFC1 value                           : 100
       Local RXFC1 value                           : 100
       Last Received SLR frame TXFCf(tc)           : 100
       Last Received SLR frame TXFCb(tc)           : 100
       SLM Frame Loss:
       Frame Loss (far-end)                        : 0 (0.00 %)
       Frame Loss (near-end)                       : 0 (0.00 %)
```

request interface link-degrade-recover

Syntax request interface link-degrade-recover *interfaces-name*

Release Information Command introduced in Junos OS Release 15.1.

Description Manually recover a degraded physical link. Manual recovery is used when the interface has any Layer 2 and Layer 3 protocols that prevents autorecovery. This command is applicable only if you have configured the manual link recovery option on the interface.

NOTE: Manual recovery option is recommended for user deployments that have static route configurations causing the remote end of the link to start forwarding packets (as soon as the physical link is up) while auto-recovery is in progress.

Options *interfaces-name*—Name of the interface.

Required Privilege Level View

Related Documentation

- Link Degrade Monitoring Overview on page 381
- link-degrade-monitor on page 1221
- thresholds on page 1395
- recovery on page 1339

List of Sample Output

Manual recovery on page 1495
Interface status when link degrade is enabled on page 1495
Interface status when the defect is active on page 1496

Output Fields When you enter this command, Junos OS displays the status of your request.

Sample Output

Manual recovery

```
user@host>run request interface link-degrade-recover xe-9/1/11
FPC 9 PIC 1 PORT 11 Link        Degrade Recovery Started
```

Interface status when link degrade is enabled

```
user@host>run show interfaces xe-9/1/11
Physical interface: xe-9/1/11, Enabled, Physical link is Up
Interface index: 181, SNMP ifIndex: 664
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
```

```
BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback:
 None, Source filtering: Disabled,
  Flow control: Enabled, Speed Configuration: Auto
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 28:8a:1c:c9:0e:32, Hardware address: 28:8a:1c:c9:0e:32
  Last flapped   : 2017-10-25 01:53:17 PDT (00:00:10 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade :
    Link Monitoring                   :  Enable
    Link Degrade Set Threshold        :  1E-8
    Link Degrade Clear Threshold      :  1E-11
    Link Degrade War Set Threshold    :  1E-9
    Link Degrade War Clear Threshold  :  1E-10
    Estimated BER                     :  <= 1E-16
    Link-degrade event                :  Seconds                Count
   State
                                         0                      0
   OK
  Interface transmit statistics: Disabled

  Logical interface xe-9/1/11.0 (Index 32368) (SNMP ifIndex 33153)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold
 cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

Interface status when the defect is active

```
user@host>run show interfaces xe-9/1/11
Physical interface: xe-9/1/11, Enabled, Physical link is Down
Interface index: 181, SNMP ifIndex: 664
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback:
 None, Source filtering: Disabled,
  Flow control: Enabled, Speed Configuration: Auto
  Pad to minimum frame size: Disabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Current address: 28:8a:1c:c9:0e:32, Hardware address: 28:8a:1c:c9:0e:32
  Last flapped   : 2017-10-25 01:54:09 PDT (00:00:03 ago)
```



```
 Input rate     : 0 bps (0 pps)
 Output rate    : 0 bps (0 pps)
 Active alarms  : LINK
 Active defects : LINK, LOCAL-FAULT
 PCS statistics                      Seconds
   Bit errors                             0
   Errored blocks                         0
 Link Degrade :
   Link Monitoring                   :  Enable
   Link Degrade Set Threshold        :  1E-8
   Link Degrade Clear Threshold      :  1E-11
   Link Degrade War Set Threshold    :  1E-9
   Link Degrade War Clear Threshold  :  1E-10
   Estimated BER                     :  1E-4
   Link-degrade event                :  Seconds              Count
  State
                                        4                    1
  Defect Active
 Interface transmit statistics: Disabled

 Logical interface xe-9/1/11.0 (Index 32368) (SNMP ifIndex 33153)
   Flags: Device-Down SNMP-Traps 0x4004000 Encapsulation: ENET2
   Input packets : 0
   Output packets: 0
   Protocol inet, MTU: 1500
   Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold
cnt: 0, NH drop cnt: 0
     Flags: Sendbcast-pkt-to-re
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
   Protocol multiservice, MTU: Unlimited
     Flags: Is-Primary
```

request interface mc-ae switchover (Multichassis Link Aggregation)

Syntax

request interface mc-ae switchover
<immediate> mcae-id *mcae-id*;
mcae-id *mcae-id*;

Release Information

Command introduced in Junos OS Release 13.3.

Description

Manually revert egress traffic from the active node to the designated preferred node of a multichassis aggregated Ethernet interface. You can use this command to manually switch over traffic to the preferred node when the **switchover-mode** statement for the multichassis aggregated Ethernet interface is configured as **non-revertive** at the **[edit interfaces ae*X* mc-ae]** hierarchy level.

NOTE: To run this command successfully, the **status-control** statement should be configured as **active** at the **[edit interfaces ae*X* mc-ae]** hierarchy level.

Options

immediate—(Optional) Trigger immediate switchover to the preferred node. If this option is not configured, Junos OS waits for the timer configured using the **revert-time** statement at the **[edit interfaces ae*X* mc-ae]** hierarchy level to expire before it triggers the switchover.

mcae-id ***mcae-id***—Triggers switchover for the specified mc-ae interface.

Required Privilege Level

view

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces*

List of Sample Output

request interface mc-ae switchover immediate mcae-id on page 1498
request interface mc-ae switchover mcae-id on page 1499

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface mc-ae switchover immediate mcae-id

```
user@host >request interface mc-ae switchover immediate mcae-id 2
MCAE: Switchover Done
```

Sample Output

request interface mc-ae switchover mcae-id

```
user@host >request interface mc-ae switchover mcae-id 2
Switchover In Progress: Please check after 1 minutes,

Use “show interfaces mc-ae revertive-info” to check for the status
```


request interface (revert | switchover) (Aggregated Ethernet Link Protection)

Syntax request interface (revert | switchover) ae*x*

Release Information Command introduced in Junos OS Release 8.3.

Description Manually revert egress traffic from the designated backup link to the designated primary link of an aggregated Ethernet interface for which link protection is enabled, or manually switch egress traffic from the primary link to the backup link. This traffic includes transit traffic and local traffic originated on the router itself.

NOTE: When link protection is enabled on an aggregated Ethernet interface, if the primary link fails, the router automatically routes egress traffic to the backup link. However, the router does not automatically route egress traffic back to the primary link when the primary link is subsequently reestablished. Instead, you manually control when to have traffic diverted back to the primary link by issuing the **request interface (revert | switchover)** (Aggregated Ethernet Link Protection) operational command and specifying the **revert** keyword.

On M Series and T Series routers, use the **request interface (revert | switchover)** (Adaptive Services) operational command to manually revert to the primary adaptive services interface or link services interface, or to switch from the primary to the secondary interface. For information about this command, see *request interface (revert | switchover) (Adaptive Services)*.

Options **revert**—Restores egress traffic processing to the primary link.

switchover—Transfers egress traffic processing to the secondary (backup) link.

ae*x*—Aggregated Ethernet logical interface number: 0 through 15.

Required Privilege Level view

List of Sample Output request interface revert on page 1500

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

request interface revert

```
user@host >request interface revert ae1
```


request lacp link-switchover

Syntax request lacp link-switchover ae*x*

Release Information Command introduced in Junos OS Release 9.3.

Description Manually switch aggregated Ethernet active or standby LACP links.

NOTE: Because this command overrides LACP priority calculations, we strongly recommend that you use this command only when the actor (in this case, the Juniper Networks router) is controlling the active or standby link and the partner (peer) is following. This scenario occurs when you configure only the actor for link protection.

Options **ae*x***—Aggregated Ethernet logical interface number: 0 through 15.

Required Privilege Level view

List of Sample Output request lacp link-switchover aeX on page 1501

Output Fields When you enter this command, you are provided feedback on the status of your request. To view the switchover, use the **show lacp interfaces** command.

Sample Output

request lacp link-switchover aeX

```
user@host >request lacp link-switchover ae0ae0: Request succeeded
```


show chassis hardware

List of Syntax

Syntax

show chassis hardware
<detail | extensive>
<clei-models>
<models>

Syntax (EX Series)

show chassis hardware
<clei-models>
<detail | extensive>
<models>
<satellite [slot-id *slot-id* |device-alias *alias-name*]>

Syntax (T4000 Router)

show chassis hardware
<clei-models>
<detail | extensive>
<models>

Syntax (TX Matrix Router)

show chassis hardware
<clei-models>
<detail | extensive>
<models>
<lcc *number* | scc>

Syntax (TX Matrix Plus Router)

show chassis hardware
<clei-models>
<detail | extensive>
<models>
<lcc *number* | sfc *number*>

Syntax (MX Series Routers)

show chassis hardware
<detail | extensive>
<clei-models>
<models>

<all-members>
<local>
<member *member-id*>

Syntax (MX104, MX204, MX2010, MX2020, MX10003, MX10008, and MX2008 3D Universal Edge Routers)

show chassis hardware
<clei-models>
<detail | extensive>
<models>
<satellite [slot-id *slot-id* |device-alias *alias-name*]>

Syntax (QFX Series)

show chassis hardware
<detail | extensive>
<clei-models>
<interconnect-device *name*>
<node-device *name*>
<models>

Syntax (OCX Series)

show chassis hardware
<detail | extensive>
<clei-models>
<models>

Syntax (PTX Series Packet Transport Routers)

show chassis hardware
<detail | extensive>
<clei-models>
<models>

Syntax (ACX Series Universal Metro Routers)

show chassis hardware
<detail | extensive>
<clei-models>
<models>

Syntax (ACX5048 and ACX5096 Routers)

show chassis hardware
<detail | extensive>
<clei-models>
<models>

Syntax (ACX500 Routers)

show chassis hardware
<detail | extensive>
<clei-models>
<models>

Release Information

Command introduced before Junos OS Release 7.4.
models option introduced in Junos OS Release 8.2.
Command introduced in Junos OS Release 9.0 for EX Series switches.
sfc option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.
Command introduced in Junos OS Release 11.1 for QFX Series.

Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 3D Universal Edge Routers.
Information for **disk** and **usb** introduced in Junos OS Release 15.1X53-D60 for QFX10002, QFX10008, and QFX10016 switches.
Command introduced in Junos OS Release 15.1X54-D20 for ACX5048 and ACX5096 Routers.
Command introduced in Junos OS Release 17.2 for MX2008 3D Universal Edge Routers.
Command introduced in Junos OS Release 17.2 for PTX10008 Routers.
Command introduced in Junos OS Release 17.3 for MX10003 3D Universal Edge Routers.
Command introduced in Junos OS Release 17.3 for MX150 Router Appliance.
Command introduced in Junos OS Release 17.4 for MX204 Routers.
Command introduced in Junos OS Release 18.1R1 for EX9251 Switches.
Command introduced in Junos OS Release 18.2 for EX9253 Switches.
Command introduced in Junos OS Release 18.2R1 for MX10008 Routers

Description Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.

In the EX Series switch command output, FPC refers to the following:

- On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC *number* is always 0.
- On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC *number* equals the member ID, from 0 through 9.
- On EX8208 and EX8216 switches—Refers to a line card; FPC *number* equals the slot number for the line card.

On QFX3500, QFX5100, and OCX Series standalone switches, and PTX1000 routers both the FPC and FPC *number* are always **0**.

On T4000 Type 5 FPCs, there are no **top temperature sensor** or **bottom temperature sensor** parameters. Instead, **fan intake temperature sensor** and **fan exhaust temperature sensors** parameters are displayed.

Starting from Junos OS Release 11.4, the output of the **show chassis hardware models** operational mode command displays the enhanced midplanes FRU model numbers (CHAS-BP3-MX240-S, CHAS-BP3-MX480-S or CHAS-BP3-MX960-S) based on the router. Prior to release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through the Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 14.1, the output of the **show chassis hardware detail | extensive | clei-models | models** operational mode command displays the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router.

Options

none—Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached routers.

clei-models—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).

detail—(Optional) Include RAM and disk information in output.

extensive—(Optional) Display ID EEPROM information.

all-members—(MX Series routers only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.

lcc *number*—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

member *member-id*—(MX Series routers and EX Series switches) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value 0 or 1.

models—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

node-device *name*—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.

satellite [slot-id *slot-id* | device-alias *alias-name*]—(Junos Fusion only) (Optional) Display hardware information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.

scc—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with **0**.

Additional Information

The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines, as shown in Table 107 on page 1506.

Table 107: Routing Engines Displaying DIMM Information

Routing Engines	Routers
RE-S-1800x2 and RE-S-1800x4	MX240, MX480, and MX960 routers
RE-A-1800x2	M120 and M320 routers

In Junos OS Release 11.4 and later, the output for the **show chassis hardware models** operational mode command for MX Series routers display the enhanced midplanes FRU model numbers—CHAS-BP3-MX240-S, CHAS-BP3-MX480-S, or CHAS-BP3-MX960-S—based on the router. In releases before Junos OS Release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 17.3R1, the output of the **show chassis hardware** command displays the mode in which vMX is running (performance mode or lite mode) in the part number field for the FPC. **RIOT-PERF** indicates performance mode and **RIOT-LITE** indicates lite mode.

Required Privilege Level

view

Related Documentation

- *show chassis power*

List of Sample Output

Output Fields Table 108 on page 1511 lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 108: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
Item	Chassis component: • (EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches. • (MX Series routers and EX Series switches)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two "pseudo" FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB). MX104 routers have a built-in Packet forwarding Engine and a Forwarding Engine Board (FEB). The Packet Forwarding Engine of the MX104 router has three "pseudo" FPCs (FPC0, FPC1, and FPC2). • (M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller. • (M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB. • (QFX Series)—Information about the chassis, Pseudo CB, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). • (PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal). • (MX2010, MX2020, and MX2008 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays. • (vMX routers)—Information about the chassis, midplane, Routing Engines, and Control Boards (CBs). Also displays information about Flexible PIC Concentrators (FPCs) and associated Modular Interface Cards (MICs) and Physical Interface Cards (PICs).	All levels
Version	Revision level of the chassis component.	All levels
Part number	Part number of the chassis component.	All levels

Table 108: show chassis hardware Output Fields (continued)

Field Name	Field Description	Level of Output
Serial number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.	All levels
Assb ID or **Assembly ID**	(**extensive** keyword only) Identification number that describes the FRU hardware.	**extensive**
Assembly Version	(**extensive** keyword only) Version number of the FRU hardware.	**extensive**
Assembly Flags	(**extensive** keyword only) Flags.	**extensive**
FRU model number	(**clei-models**, **extensive**, and **models** keyword only) Model number of the FRU hardware component.	none specified
CLEI code	(**clei-models** and **extensive** keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
EEPROM Version	ID EEPROM version used by the hardware component: **0x00** (version 0), **0x01** (version 1), or **0x02** (version 2).	**extensive**
Description	Brief description of the hardware item: • Type of power supply. • Type of PIC. If the PIC type is not supported on the current software release, the output states **Hardware Not Supported**. • Type of FPC: **FPC Type 1**, **FPC Type 2**, **FPC Type 3**, **FPC Type 4** , or **FPC TypeOC192**. On EX Series switches, a brief description of the FPC. The following list shows the PIM abbreviation in the output and the corresponding PIM name. ◦ **2x FE**—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM ◦ **4x FE**—4-port Fast Ethernet ePIM ◦ **1x GE Copper**—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port) ◦ **1x GE SFP**—SFP Gigabit Ethernet ePIM (one fiber port) ◦ **2x Serial**—Dual-port serial PIM ◦ **2x T1**—Dual-port T1 PIM ◦ **2x E1**—Dual-port E1 PIM ◦ **2x CT1E1**—Dual-port channelized T1/E1 PIM ◦ **1x T3**—T3 PIM (one port) ◦ **1x E3**—E3 PIM (one port) ◦ **4x BRI S/T**—4-port ISDN BRI S/T PIM ◦ **4x BRI U**—4-port ISDN BRI U PIM ◦ **1x ADSL Annex A**—ADSL 2/2+ Annex A PIM (one port, for POTS)	All levels

Table 108: show chassis hardware Output Fields (continued)

Field Name	Field Description	Level of Output
	- **1x ADSL Annex B**—ADSL 2/2+ Annex B PIM (one port, for ISDN) - **2x SHDSL (ATM)**—G SHDSL PIM (2-port two-wire module or 1-port four-wire module) - **1x TGM550**—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog **LINE** ports, and two analog **TRUNK** ports) - **1x DS1 TIM510**—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup) - **4x FXS, 4xFX0, TIM514**—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog **LINE** ports and four analog **TRUNK** ports) - **4x BRI TIM521**—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports) - **Crypto Accelerator Module**—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services • **MPC M 16x 10GE**—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.) • For hosts, the Routing Engine type. • For small form-factor pluggable transceiver (SFP) modules, the type of fiber: **LX**, **SX**, **LH**, or **T**. • LCD description for EX Series switches (except EX2200 switches). • **MPC2**—1-port MPC2 that supports two separate slots for MICs. • **MPC3E**—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs. • 100GBASE-LR4, pluggable CFP optics • Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy. • Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs). • **MPC4E**—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE on MX2020, MX960, MX480, and MX240 routers. • LCD description for MX Series routers	

Sample Output

show chassis hardware (EX8216 Switch)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis          REV 06                CY0109220035      EX8216
Midplane         REV 06   710-016845   BA0909120112      EX8216-MP
CB 0             REV 22   710-020771   AX0109197723      EX8216-RE320
CB 1             REV 22   710-020771   AX0109197726      EX8216-RE320
  Routing Engine 1        BUILTIN      BUILTIN           RE-EX8216
```



```
FPC 3            REV 19   710-020683   BC0109083125      EX8200-48F
  CPU            REV 13   710-020598   BF0109144549      EX8200-CPU
FPC 4            REV 17   710-020683   BC0108500127      EX8200-48F
  CPU            REV 10   710-020598   BF0108460510      EX8200-CPU
  PIC 0                   BUILTIN      BUILTIN           48x 100 Base-QFX/1000
Base-X
    Xcvr 1       REV 01   740-011613   PE70V89           SFP-SX
    Xcvr 11      REV 01   740-011613   PE70YCE           SFP-SX
    Xcvr 12      REV 01   740-011613   PE70VSH           SFP-SX
    Xcvr 13      REV 01   740-011613   E08C02063         SFP-SX
    Xcvr 14      REV 01   740-011613   PE70VKU           SFP-SX
    Xcvr 15      REV 01   740-011613   E08E03372         SFP-SX
    Xcvr 21      REV 01   740-011613   PE70VAD           SFP-SX
    Xcvr 22      REV 01   740-011613   E08E01228         SFP-SX
    Xcvr 23      REV 01   740-011613   PE70VSL           SFP-SX
    Xcvr 24      REV 01   740-011613   E08E03409         SFP-SX
    Xcvr 25      REV 01   740-011613   PE70VL4           SFP-SX
    Xcvr 26      REV 01   740-011613   PDQ4L2Z           SFP-SX
    Xcvr 27      REV 01   740-011613   PE70WFK           SFP-SX
    Xcvr 28      REV 01   740-011782   PBD2B5U           SFP-SX
    Xcvr 29      REV 01   740-011613   PE70UQX           SFP-SX
    Xcvr 30      REV 01   740-011613   PE70VL5           SFP-SX
    Xcvr 31      REV 01   740-011613   PE70V0F           SFP-SX
    Xcvr 32      REV 01   740-011613   E08C02052         SFP-SX
    Xcvr 33      REV 01   740-011613   E08C02197         SFP-SX
    Xcvr 34      REV 01   740-011613   PE70V0L           SFP-SX
    Xcvr 35      REV 01   740-011613   E08E03390         SFP-SX
    Xcvr 36      REV 01   740-011613   PDQ4VL9           SFP-SX
    Xcvr 37      REV 01   740-011613   E08E03370         SFP-SX
    Xcvr 38      REV 01   740-011613   E08E03362         SFP-SX
    Xcvr 39      REV 01   740-011613   E08C02065         SFP-SX
    Xcvr 40      REV 01   740-011613   E08E03405         SFP-SX
    Xcvr 41      REV 01   740-011613   E08E03411         SFP-SX
    Xcvr 43      REV 01   740-011613   E08C02171         SFP-SX
    Xcvr 45      REV 01   740-011613   E08E03410         SFP-SX
FPC 13           REV 16   710-016837   BB0109051344      EX8200-8XS
  CPU
SIB 0            REV 10   710-021613   AY0109166244      EX8216-SF320
SIB 1            REV 10   710-021613   AY0109166357      EX8216-SF320
SIB 2            REV 10   710-021613   AY0109166362      EX8216-SF320
SIB 3            REV 10   710-021613   AY0109166338      EX8216-SF320
SIB 4            REV 10   710-021613   AY0109166350      EX8216-SF320
SIB 5            REV 10   710-021613   AY0109166365      EX8216-SF320
SIB 6            REV 10   710-021613   AY0109166361      EX8216-SF320
SIB 7            REV 10   710-021613   AY0109166399      EX8216-SF320
PSU 0            REV 17   740-021466   BG0709170003      EX8200-AC2K
PSU 1            REV 17   740-021466   BG0709170004      EX8200-AC2K
PSU 2            REV 17   740-021466   BG0709170020      EX8200-AC2K
PSU 3            REV 17   740-021466   BG0709170017      EX8200-AC2K
PSU 4            REV 17   740-021466   BG0709170008      EX8200-AC2K
PSU 5            REV 17   740-021466   BG0709170018      EX8200-AC2K
Top Fan Tray
  FTC 0          REV 4    760-022620   CX1209140212      EX8216-FT
  FTC 1          REV 4    760-022620   CX1209140212      EX8216-FT
Bottom Fan Tray
  FTC 0          REV 4    760-022620   CX1209140211      EX8216-FT
  FTC 1          REV 4    760-022620   CX1209140211      EX8216-FT
LCD 0            REV 04   710-025742   CE0109186919      EX8200 LCD
```


show chassis hardware clei-models (EX8216 Switch)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 08   710-016845
PSU 0            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 1            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 2            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 3            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 4            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 5            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
Top Fan Tray
Bottom Fan Tray
```

show chassis hardware clei-models (T1600 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 03   710-005608                     CHAS-BP-T640-S
FPM Display      REV 05   710-002897                     CRAFT-T640-S
CIP              REV 06   710-002895                     CIP-L-T640-S
PEM 0            Rev 07   740-017906   IPUPAC7KTA        PWR-T1600-3-80-DC-S
PEM 1            Rev 18   740-002595                     PWR-T-DC-S
SCG 0            REV 15   710-003423                     SCG-T-S
Routing Engine 0 REV 08   740-014082                     RE-A-2000-4096-S
Routing Engine 1 REV 07   740-014082                     RE-A-2000-4096-S
CB 0             REV 05   710-007655                     CB-T-S
CB 1             REV 03   710-017707                     CB-T-S
FPC 0            REV 07   710-013558                     T640-FPC2-E2
  PIC 0          REV 01   750-010618                     PB-4GE-SFP
  PIC 1          REV 06   750-001900                     PB-1OC48-SON-SMSR
  PIC 2          REV 14   750-001901                     PB-4OC12-SON-SMIR
  PIC 3          REV 07   750-001900                     PB-1OC48-SON-SMSR
FPC 1            REV 06   710-013553                     T640-FPC1-E2
  PIC 0          REV 08   750-001072                     P-1GE-SX
  PIC 1          REV 10   750-012266                     PB-4GE-TYPE1-SFP-IQ2
  PIC 2          REV 22   750-005634                     PB-1CHOC12SMIR-QPP
FPC 2
  PIC 0          REV 16   750-007141                     PC-10GE-SFP
  PIC 1          REV 06   750-015217                     PC-8GE-TYPE3-SFP-IQ2
  PIC 2          REV 05   750-004695                     PC-TUNNEL
  PIC 3          REV 17   750-009553                     PC-4OC48-SON-SFP
FPC 3            REV 01   710-010154                     T640-FPC3-E
  PIC 0          REV 07   750-012793                     PC-1XGE-TYPE3-XFP-IQ2
  PIC 1          REV 25   750-007141                     PC-10GE-SFP
  PIC 2          REV 17   750-009553                     PC-4OC48-SON-SFP
  PIC 3          REV 32   750-003700                     PC-1OC192-SON-VSR
FPC 4            REV 16   710-013037                     T1600-FPC4-ES
 PIC 1           REV 06   750-034781                     PD-1CE-CFP
FPC 5            REV 02   710-013037                     T1600-FPC4-ES
  PIC 0          REV 16   750-012518                     PD-4OC192-SON-XFP
  PIC 1          REV 01   750-010850                     PD-1OC768-SON-SR
FPC 6            REV 14   710-013037                     T1600-FPC4-ES
  PIC 0          REV 11   750-017405                     PD-4XGE-XFP
  PIC 1          REV 13   750-017405                     PD-4XGE-XFP
FPC 7            REV 09   710-007529                     T640-FPC3
  PIC 0          REV 10   750-012793                     PC-1XGE-TYPE3-XFP-IQ2
  PIC 1          REV 01   750-015217                     PC-8GE-TYPE3-SFP-IQ2
```

```
  PIC 2          REV 01    750-015217                    PC-8GE-TYPE3-SFP-IQ2
  PIC 3          REV 15    750-009450                    PC-10C192-SON-SR2
SIB 0            REV 07    710-013074                    SIB-I-T1600-S
SIB 1            REV 07    710-013074                    SIB-I-T1600-S
SIB 2            REV 07    710-013074                    SIB-I-T1600-S
SIB 3            REV 07    710-013074                    SIB-I-T1600-S
SIB 4            REV 07    710-013074                    SIB-I-T1600-S
Fan Tray 0                                               FANTRAY-T-S
Fan Tray 1                                               FANTRAY-T-S
Fan Tray 2                                               FAN-REAR-TX-T640-S
```

show chassis hardware clei-models (PTX10008 Routers)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 27   750-054097   CMMUM00ARA        QFX10008-CHAS
CB 0             REV 02   750-068820   CMUCAH3CTB        QFX10000-RE
CB 1             REV 02   750-068820   CMUCAH3CTB        QFX10000-RE
FPC 0            REV 36   750-051354   CMUIAM9BAA        QFX10000-36Q
  PIC 0                   BUILTIN
FPC 1            REV 33   750-051354   CMUIAM9BAA        QFX10000-36Q
  PIC 0                   BUILTIN
FPC 2            REV 32   750-051357   CMUIANABAA        QFX10000-30C
  PIC 0                   BUILTIN
FPC 3            REV 35   750-051357   CMUIANABAA        QFX10000-30C
  PIC 0                   BUILTIN
FPC 5            REV 08   750-068822   CMUIAM9BAB        QFX10000-36Q
  PIC 0                   BUILTIN
FPC 6            REV 08   750-068822   CMUIAM9BAB        QFX10000-36Q
  PIC 0                   BUILTIN
FPD Board        REV 07   711-054687
Power Supply 0   REV 02   740-049388   CMUPADNBAA        QFX10000-PWR-AC
Power Supply 1   REV 02   740-049388   CMUPADNBAA        QFX10000-PWR-AC
Power Supply 2   REV 02   740-049388   CMUPADNBAA        QFX10000-PWR-AC
Power Supply 3   REV 02   740-049388   CMUPADNBAA        QFX10000-PWR-AC
Power Supply 4   REV 02   740-049388   CMUPADNBAA        QFX10000-PWR-AC
Power Supply 5   REV 02   740-049388   CMUPADNBAA        QFX10000-PWR-AC
FTC 0            REV 14   750-050108   CMUCAHZCAA        QFX10008-FAN-CTRL
FTC 1            REV 14   750-050108   CMUCAHZCAA        QFX10008-FAN-CTRL
Fan Tray 0       REV 09   760-054372   CMUCAHYCAA        QFX10008-FAN
Fan Tray 1       REV 09   760-054372   CMUCAHYCAA        QFX10008-FAN
SIB 0            REV 24   750-050058   CMUCAH0CAA        QFX10008-SF
SIB 1            REV 24   750-050058   CMUCAH0CAA        QFX10008-SF
SIB 2            REV 24   750-050058   CMUCAH0CAA        QFX10008-SF
SIB 3            REV 24   750-050058   CMUCAH0CAA        QFX10008-SF
SIB 4            REV 24   750-050058   CMUCAH0CAA        QFX10008-SF
SIB 5            REV 23   750-050058   CMUCAH0CAA        QFX10008-SF
```

show chassis hardware clei-models (PTX10016 Routers)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 24   750-077138   CMMUN00ARA        JNP10016
CB 0             REV 04   711-065897   PROTOXCLEI        PROTO-ASSEMBLY
CB 1             REV 05   711-065897   PROTOXCLEI        PROTO-ASSEMBLY
FPC 2
  PIC 0                   BUILTIN
FPC 4            REV 35   750-071976   CMUIANABAA        JNP10K-LC1101
```



```
  PIC 0                      BUILTIN
FPC 5            REV 13     750-068822    CMUIAM9BAC         QFX10000-36Q
  PIC 0                      BUILTIN
FPC 6            REV 41     750-071976    CMUIANABAB         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 7            REV 35     750-071976    CMUIANABAA         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 8            REV 35     750-071976    CMUIANABAA         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 9            REV 41     750-071976    CMUIANABAB         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 10           REV 35     750-071976    CMUIANABAA         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 11           REV 35     750-071976    CMUIANABAA         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 13           REV 41     750-071976    CMUIANABAB         JNP10K-LC1101
  PIC 0                      BUILTIN
FPC 15           REV 37     750-071976    CMUIANABAA         JNP10K-LC1101
  PIC 0                      BUILTIN
Power Supply 0   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 1   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 2   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 3   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 4   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 5   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 6   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 7   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 8   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Power Supply 9   REV 01     740-073147    CMUPADPBAA         JNP10K-PWR-DC
Fan Tray 0                                                   QFX5100-FAN-AFO
Fan Tray 1                                                   QFX5100-FAN-AFO
SIB 0            REV 15     750-077140    CMUCAH6CAA         JNP10016-SF
SIB 1            REV 15     750-077140    CMUCAH6CAA         JNP10016-SF
SIB 2            REV 15     750-077140    CMUCAH6CAA         JNP10016-SF
SIB 3            REV 15     750-077140    CMUCAH6CAA         JNP10016-SF
SIB 4            REV 15     750-077140    CMUCAH6CAA         JNP10016-SF
SIB 5            REV 15     750-077140    CMUCAH6CAA         JNP10016-SF
FPD Board        REV 07     711-054687
```

show chassis hardware (EX2300-C Switch)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                HV0215410003      EX2300-C-12P
Pseudo CB 0
Routing Engine 0            BUILTIN    BUILTIN           RE-EX2300C-12P
FPC 0            REV 04     650-059984 HV0215410003      EX2300-C-12P
  CPU                       BUILTIN    BUILTIN           FPC CPU
  PIC 0          REV 04     BUILTIN    BUILTIN           12x10/100/1000 Base-T
  PIC 1          REV 04     650-059984 HV0215410003      2x10G SFP/SFP+
    Xcvr 0       REV 01     740-021309 T09K00695         SFP+-10G-LR
    Xcvr 1       REV 01     740-030658 AD1146A05JT       SFP+-10G-USR
Power Supply 0                                           JPSU-170W-AC
```

show chassis hardware (EX2300 Switch)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
```

```
Chassis                                JY0215410033      EX2300-24P
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           RE-EX2300-24P
FPC 0            REV 05   650-059968   JY0215410033      EX2300-24P
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0          REV 05   BUILTIN      BUILTIN           24x10/100/1000 Base-T
  PIC 1          REV 05   650-059968   JY0215410033      4x10G SFP/SFP+
    Xcvr 0       REV 01   740-030658   AD1125A03ES       SFP+-10G-USR
    Xcvr 1       REV 01   740-021308   AJP0TDZ           SFP+-10G-SR
    Xcvr 3       REV 01   740-021309   A9401FL           SFP+-10G-LR
Power Supply 0                                           JPSU-450W-AC-AFO
Fan Tray 0                                              Fan Module, Airflow Out
 (AFO)
Fan Tray 1                                              Fan Module, Airflow Out
 (AFO)
```

show chassis hardware detail (EX4200 Switch)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BM0208327733      EX4200-24T
Routing Engine 0 REV 11   750-021256   BM0208327733      EX4200-24T, 8 POE
Routing Engine 0                       BM0208327733      EX4200-24T, 8 POE
FPC 0            REV 11   750-021256   BM0208327733      EX4200-24T, 8 POE
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           24x 10/100/1000 Base-T
  PIC 1          REV 03B  711-021270   AR0208162285      4x GE SFP
  BRD            REV 08   711-021264   AK0208328289      EX4200-24T, 8 POE
Power Supply 0   REV 03   740-020957   AT0508346354      PS 320W AC
Fan Tray                                                 Fan Tray
```

show chassis hardware (EX4300 Switch)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                PD3713160055      EX4300-48P
Routing Engine 0 REV 04   650-044930   PD3713160055      EX4300-48P
FPC 0            REV 04   650-044930   PD3713160055      EX4300-48P
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0          REV 04   BUILTIN      BUILTIN           48x 10/100/1000 Base-T
  PIC 1          REV 04   BUILTIN      BUILTIN           4x 40GE
Power Supply 0   REV 01   740-046871   1EDA3090026       JPSU-1100-AC-AFO-A
Fan Tray 0                                              Fan Module, Airflow Out
 (AFO)
Fan Tray 1                                              Fan Module, Airflow Out
 (AFO)
```

show chassis hardware models (EX4500 Switch)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Routing Engine 0 REV 01   750-035700   GG0210271867      EX4500-40F-FB-C
FPC 0            REV 01   750-035700   GG0210271867      EX4500-40F-FB-C
  PIC 0                   BUILTIN      BUILTIN           EX4500-40F-FB-C
Power Supply 1   REV 01   740-029654   H884FS00JC09      EX4500-PWR1-AC-FB
```


show chassis hardware detail (EX9200 Switch)

user@switch> **show chassis hardware**

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN111DA44RFB      EX9208
Midplane         REV 05   710-017414   TS2912            EX9208-BP
FPM Board        REV 02   710-017254   XN1804            Front Panel Display
PEM 0            Rev 01   740-022697   QCS0906C033       PS 1.2-1.7kW; 100-240V
AC in
PEM 1            Rev 01   740-022697   QCS0906C095       PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 08   740-031116   9009122883        RE-S-EX9200-1800X4
CB 0             REV 16   750-031391   CAAW4391          EX9200-SCBEF
PC 0             REV 07   750-049612   CABJ9312          EX9200 40x1G Copper
  CPU          REV 04   711-038484   CABH8268          MPCE PMB 2G
  MIC 0          REV 02   750-049607   CABT9623          40x 1GE RJ45
    PIC 0                 BUILTIN      BUILTIN           10x 1GE RJ45
    PIC 1                 BUILTIN      BUILTIN           10x 1GE RJ45
    PIC 2                 BUILTIN      BUILTIN           10x 1GE RJ45
    PIC 3                 BUILTIN      BUILTIN           10x 1GE RJ45
FPC 1            REV 10   710-013699   CAAN3529          EX9200-40x1G-SFP
  CPU            REV 04   711-038484   CAAL7608          MPCE PMB 2G
   MIC 0           REV 26   750-028392   CAAS5151          20x 1GE SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE SFP
    PIC 1                 BUILTIN      BUILTIN           10x 1GE SFP
  MIC 1          REV 26   750-028392   CAAC8006          20x 1GE SFP
    PIC 2                 BUILTIN      BUILTIN           10x 1GE SFP
      Xcvr 8     REV 01   740-011613   E08L03674         SFP-SX
      Xcvr 9     REV 01   740-011613   E08M00243         SFP-SX
    PIC 3                 BUILTIN      BUILTIN           10x 1GE SFP
    FPC 3        REV 10   710-013699   CAAR5261          EX9200-40x1G-SFP
  CPU            REV 04   711-038484   CAAS2118          MPCE PMB 2G
   MIC 0           REV 26   750-028392   CAAS5067          20x 1GE SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE SFP
      Xcvr 2     REV 01   740-031851   PNA7L8U           SFP-SX
      Xcvr 3     REV 02   740-011613   AM0943SEKGZ       SFP-SX
      Xcvr 4     REV 02   740-011613   AM0943SEJZ9       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE SFP
  MIC 1          REV 26   750-028392   CAAS5132          20x 1GE SFP
    PIC 2                 BUILTIN      BUILTIN           10x 1GE SFP
      Xcvr 4     REV 01   740-011613   E08D02625         SFP-SX
      Xcvr 9     REV 02   740-011613   PJH4RD9           SFP-SX
    PIC 3                 BUILTIN      BUILTIN           10x 1GE SFP
      Xcvr 0     REV 01   740-011613   AM0813S8YME       SFP-SX
      Fan Tray                                           Left Fan Tray
```

show chassis hardware detail (EX9251 Switch)

user@switch> **show chassis hardware**

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BLANK             EX9251
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x6
CB 0             REV 05   750-069579   CAGT1382          EX9251
FPC 0                     BUILTIN      BUILTIN           MPC
  PIC 0                   BUILTIN      BUILTIN           4XQSFP28 PIC
    Xcvr 0       REV 01   740-044512   APF14500007NHC    QSFP+-40G-CU50CM
    Xcvr 2       REV 01   740-046565   QH21035H          QSFP+-40G-SR4
```



```
  PIC 1                 BUILTIN      BUILTIN           8XSFPP PIC
    Xcvr 0      REV 01  740-031980   AA15393URH7       SFP+-10G-SR
    Xcvr 1      REV 01  740-031980   AA162832LVG       SFP+-10G-SR
    Xcvr 2      REV 01  740-031980   MXA0NKJ           SFP+-10G-SR
    Xcvr 3      REV 01  740-031980   MXA0K75           SFP+-10G-SR
    Xcvr 4      REV 01  740-021308   MXA138L           SFP+-10G-SR
    Xcvr 5      REV 01  740-021308   13T511102684      SFP+-10G-SR
    Xcvr 6      REV 01  740-021308   MXA138E           SFP+-10G-SR
    Xcvr 7      REV 01  740-021308   MXA152N           SFP+-10G-SR
PEM 0           REV 02  740-070749   1F186390060       AC AFO 650W PSU
PEM 1           REV 02  740-070749   1F186390045       AC AFO 650W PSU
Fan Tray 0                                             Fan Tray, Front to Back
 Airflow - AFO
Fan Tray 1                                             Fan Tray, Front to Back
 Airflow - AFO
```

show chassis hardware detail (EX9253 Switch)

user@switch> **show chassis hardware**

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN126145CJCB      EX9253
Midplane         REV 06   750-074276   CAJE4108          Midplane 2
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x6
Routing Engine 1          BUILTIN      BUILTIN           RE-S-2X00x6
CB 0             REV 24   750-067071   CAJF6414          Control Board
  Mezz           REV 14   711-066896   CAJF6327          Control Mezz Board
CB 1             REV 24   750-067071   CAJF6398          Control Board
  Mezz           REV 14   711-066896   CAJF6314          Control Mezz Board
FPC 0            REV 19   750-066879   CAJD1692          LC2103
  CPU                     BUILTIN      BUILTIN           SMPC PMB
  PIC 0                   BUILTIN      BUILTIN           6xQSFPP
    Xcvr 0       REV 01   740-054053   QH20019A          QSFP+-4X10G-SR
  PIC 1          REV 15   750-068806   CAJD1416          MIC1
    Xcvr 0       REV 01   740-061405   1ECQ1151163       QSFP-100GBASE-SR4
    Xcvr 1       REV 01   740-061405   1ECQ11511AK       QSFP-100GBASE-SR4
    Xcvr 2       REV 01   740-032986   QB160112          QSFP+-40G-SR4
FPC 1            REV 19   750-066879   CAJD1685          LC2103
  CPU                     BUILTIN      BUILTIN           SMPC PMB
  PIC 0                   BUILTIN      BUILTIN           6xQSFPP
  PIC 1          REV 15   750-068806   CAJD1393          MIC1
    Xcvr 0       REV 01   740-032986   QB120887          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-032986   QD465034          QSFP+-40G-SR4
    Xcvr 2       REV 01   740-052009   UWE2CBQ           QSFP+-40G-LR4
    Xcvr 4       REV 01   740-032986   QB120701          QSFP+-40G-SR4
PEM 0            REV 01   740-066937   1HS17070027       JNP-PWR1600-AC
PEM 1            REV 01   740-066937   1HS17070151       JNP-PWR1600-AC
PEM 4            REV 01   740-066937   1HS17070090       JNP-PWR1600-AC
PEM 5            REV 01   740-066937   1HS16480119       JNP-PWR1600-AC
Fan Tray 0       REV 08   760-069329   CAJF6944          JNP FAN 3RU
Fan Tray 1       REV 08   760-069329   CAJF6863          JNP FAN 3RU
Fan Tray 2       REV 08   760-069329   CAJF6891          JNP FAN 3RU
Fan Tray 3       REV 08   760-069329   CAJF6937          JNP FAN 3RU
```

show chassis hardware detail (PTX10008 Routers)

user@switch> **show chassis hardware detail**

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
```



```
Chassis                                  DE487              JNP10008 [PTX10008 -
PILOT BUILD V1.1]
Midplane         REV 27   750-054097   ACPD4307           Midplane 8
Routing Engine 0          BUILTIN      BUILTIN            RE-PTX-2X00x4
  vtbd0 15360 MB                                          Virtio Block Disk
  vtbd1 15360 MB                                          Virtio Block Disk
  ada0    128 MB  QEMU                 QM00002            Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0      Intel              uhub0
  usb1 (addr 0.2) product 0x0020 32    vendor 0x8087      uhub1
Routing Engine 1          BUILTIN      BUILTIN            RE-PTX-2X00x4
  vtbd0 15360 MB                                          Virtio Block Disk
  vtbd1 15360 MB                                          Virtio Block Disk
  ada0    128 MB  QEMU                 QM00002            Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0      Intel              uhub0
  usb1 (addr 0.2) product 0x0020 32    vendor 0x8087      uhub1
CB 0             REV 02   750-068820   ACNZ4440           Control Board
CB 1             REV 02   750-068820   ACNZ8284           Control Board
FPC 0            REV 36   750-051354   ACNP4679           LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN      BUILTIN            FPC CPU
  PIC 0                  BUILTIN      BUILTIN           12x100GE/36x40GE/144x10GE

    Xcvr 1       REV 01   740-058734   1ECQ113834D        QSFP-100GBASE-SR4
    Xcvr 5       REV 01   740-058734   1ECQ1137067        QSFP-100GBASE-SR4
    Xcvr 6       REV 01   740-054053   QF3205SD           QSFP+-4X10G-SR
    Xcvr 7       REV 01   740-058734   1ECQ11381MP        QSFP-100GBASE-SR4
    Xcvr 11      REV 01   740-061405   1ACQ110507K        QSFP-100GBASE-SR4
    Xcvr 13      REV 01   740-058734   1ECQ11390ZB        QSFP-100GBASE-SR4
    Xcvr 17      REV 01   740-058734   1ECQ11381M1        QSFP-100GBASE-SR4
    Xcvr 19      REV 01   740-058734   1ECQ11381JS        QSFP-100GBASE-SR4
    Xcvr 23      REV 01   740-058734   1ACQ112000E        QSFP-100GBASE-SR4
    Xcvr 25      REV 01   740-058734   1ECQ11381NT        QSFP-100GBASE-SR4
    Xcvr 28      REV 01   740-054053   QG1502WV           QSFP+-4X10G-SR
    Xcvr 29      REV 01   740-058734   1ACQ112000D        QSFP-100GBASE-SR4
    Xcvr 33      REV 01   740-058734   1ACQ1134065        QSFP-100GBASE-SR4
    Xcvr 34      REV 01   740-067442   XV20L4L            QSFP+-40G-SR4
FPC 1            REV 33   750-051354   ACNX8831           LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN      BUILTIN            FPC CPU
  PIC 0                  BUILTIN      BUILTIN           12x100GE/36x40GE/144x10GE

    Xcvr 5                NON-JNPR     37700171YY0084     QSFP-100GBASE-LR4
    Xcvr 25               NON-JNPR     GDA2017459         QSFP-100GBASE-LR4
    Xcvr 29               NON-JNPR     GDF2008750         QSFP-100GBASE-LR4
FPC 2           REV 32   750-051357   ACPB0341          LC1101 - 30C / 30Q / 96X

  CPU                     BUILTIN      BUILTIN            FPC CPU
  PIC 0                   BUILTIN      BUILTIN           30x100GE/30x40GE/96x10GE

    Xcvr 0                NON-JNPR     37700170YZC305     QSFP-100GBASE-LR4
    Xcvr 4                NON-JNPR     37700170YZC306     QSFP-100GBASE-LR4
    Xcvr 9       REV 01   740-054053   QF36013S           QSFP+-4X10G-SR
    Xcvr 12      REV 01   740-067442   XV301AU            QSFP+-40G-SR4
    Xcvr 14      REV 01   740-043308   UWE2CG9            QSFP+-40G-LR4
    Xcvr 16      REV 01   740-043308   UWH141S            QSFP+-40G-LR4
    Xcvr 17      REV 01   740-058734   1ECQ11180VH        QSFP-100GBASE-SR4
    Xcvr 18      REV 01   740-054050   INFAJ0492237       QSFP+-4X10G-LR
    Xcvr 26      REV 01   740-058734   1ACQ111803N        QSFP-100GBASE-SR4
    Xcvr 27      REV 01   740-058734   1ACQ113405S        QSFP-100GBASE-SR4
FPC 3           REV 35   750-051357   ACPD2186          LC1101 - 30C / 30Q / 96X
```



```
  CPU                    BUILTIN      BUILTIN          FPC CPU
  PIC 0                 BUILTIN      BUILTIN         30x100GE/30x40GE/96x10GE

    Xcvr 0        REV 01   740-061409   1GCQA1470A3      QSFP-100GBASE-LR4
    Xcvr 1        REV 01   740-061409   1GCQA1470XC      QSFP-100GBASE-LR4
    Xcvr 7                 NON-JNPR     FG4550500008     QSFP-100G-CWDM4
    Xcvr 24       REV 01   740-058734   1ECQ11381LX      QSFP-100GBASE-SR4
    Xcvr 29       REV 01   740-043308   UWE0UYS          QSFP+-40G-LR4
FPC 5             REV 08   750-068822   ACPF0057         LC1102 - 12C / 36Q /
144X
  CPU                    BUILTIN      BUILTIN          FPC CPU
  PIC 0                 BUILTIN      BUILTIN         12x100GE/36x40GE/144x10GE
FPC 6             REV 08   750-068822   ACPE9951         LC1102 - 12C / 36Q /
144X
  CPU                    BUILTIN      BUILTIN          FPC CPU
  PIC 0                 BUILTIN      BUILTIN         12x100GE/36x40GE/144x10GE

    Xcvr 1        REV 01   740-054053   QF3208LG         QSFP+-4X10G-SR
    Xcvr 7        REV 01   740-067442   XV20LGN          QSFP+-40G-SR4
    Xcvr 8        REV 01   740-067442   XV20VMV          QSFP+-40G-SR4
    Xcvr 9        REV 01   740-067442   XV20KCN          QSFP+-40G-SR4
    Xcvr 10       REV 01   740-067442   XU504QD          QSFP+-40G-SR4
    Xcvr 11       REV 01   740-067442   XU504X7          QSFP+-40G-SR4
    Xcvr 12       REV 01   740-067442   XU504W8          QSFP+-40G-SR4
    Xcvr 16       REV 01   740-032986   QF4301JP         QSFP+-40G-SR4
    Xcvr 17       REV 01   740-032986   QF4303AE         QSFP+-40G-SR4
    Xcvr 18       REV 01   740-054050   INFAJ0492400     QSFP+-4X10G-LR
    Xcvr 19       REV 01   740-054050   INFAJ0492142     QSFP+-4X10G-LR
    Xcvr 24       REV 01   740-032986   QF4301KB         QSFP+-40G-SR4
    Xcvr 25       REV 01   740-032986   QF4303YP         QSFP+-40G-SR4
    Xcvr 30       REV 01   740-067442   XV300ZX          QSFP+-40G-SR4
    Xcvr 31       REV 01   740-043308   UWH2KBW          QSFP+-40G-LR4
    Xcvr 34       REV 01   740-054053   QG1501YU         QSFP+-4X10G-SR
FPD Board         REV 07   711-054687   ACPC7142         Front Panel Display
Power Supply 0    REV 02   740-049388   1EDL62102N9      Power Supply AC
Power Supply 1    REV 02   740-049388   1EDL60300KX      Power Supply AC
Power Supply 2    REV 02   740-049388   1EDL60300DL      Power Supply AC
Power Supply 3    REV 02   740-049388   1EDL61701BT      Power Supply AC
Power Supply 4    REV 02   740-049388   1EDL62102P7      Power Supply AC
Power Supply 5    REV 02   740-049388   1EDL62102PP      Power Supply AC
FTC 0             REV 14   750-050108   ACPE4038         Fan Controller 8
FTC 1             REV 14   750-050108   ACPE4032         Fan Controller 8
Fan Tray 0        REV 09   760-054372   ACPD6799         Fan Tray 8
Fan Tray 1        REV 09   760-054372   ACNZ3584         Fan Tray 8
SIB 0             REV 24   750-050058   ACPD4587         Switch Fabric 8
SIB 1             REV 24   750-050058   ACNZ0635         Switch Fabric 8
SIB 2             REV 24   750-050058   ACPD4908         Switch Fabric 8
SIB 3             REV 24   750-050058   ACNZ0617         Switch Fabric 8
SIB 4             REV 24   750-050058   ACNZ0527         Switch Fabric 8
SIB 5             REV 23   750-050058   ACNX6980         Switch Fabric 8
```

show chassis hardware detail (PTX10016 Routers)

```
user@switch> show chassis hardware detail

Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                 DH995            JNP10016 [PTX10016]
Midplane          REV 22   750-056555   ACPM7810         Midplane 16
Routing Engine 0           BUILTIN      BUILTIN          RE-PTX-2X00x4
  vtbd0 15360 MB                                         Virtio Block Disk
```

```
  vtbd1 15360 MB                                        Virtio Block Disk
  ada0    128 MB  QEMU                   QM00002        Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0        Intel          uhub0
  usb1 (addr 0.2) product 0x0020 32      vendor 0x8087  uhub1
Routing Engine 1          BUILTIN        BUILTIN        RE-PTX-2X00x4
  vtbd0 15360 MB                                        Virtio Block Disk
  vtbd1 15360 MB                                        Virtio Block Disk
  ada0    128 MB  QEMU                   QM00002        Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0        Intel          uhub0
  usb1 (addr 0.2) product 0x0020 32      vendor 0x8087  uhub1
CB 0             REV 03   750-068820     ACPL7238       Control Board
CB 1             REV 03   750-068820     ACPL7298       Control Board
FPC 1            REV 36   750-077140     ACNP4590       LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN        BUILTIN        FPC CPU
  PIC 0                  BUILTIN        BUILTIN        12x100GE/36x40GE/144x10GE

    Xcvr 0       REV 01   740-054053     QF3600AV       QSFP+-4X10G-SR
    Xcvr 35      REV 01   740-061405     1ACQ110507K    QSFP-100GBASE-SR4
FPC 3            REV 07   750-071975     CAHA2224       LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN        BUILTIN        FPC CPU
  PIC 0                  BUILTIN        BUILTIN        12x100GE/36x40GE/144x10GE

    Xcvr 0       REV 01   740-054053     QG1505YM       QSFP+-4X10G-SR
    Xcvr 11               NON-JNPR       GDA2017459     QSFP-100GBASE-LR4
    Xcvr 35               NON-JNPR       GDF2008750     QSFP-100GBASE-LR4
FPC 5            REV 13   750-068822     ACPD6501       LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN        BUILTIN        FPC CPU
  PIC 0                  BUILTIN        BUILTIN        12x100GE/36x40GE/144x10GE

    Xcvr 1       REV 01   740-058734     1ECQ11381LA    QSFP-100GBASE-SR4
    Xcvr 2       REV 01   740-043308     UWH141S        QSFP+-40G-LR4
    Xcvr 3       REV 01   740-043308     UWE2CG9        QSFP+-40G-LR4
FPC 6            REV 37   750-077140     ACNS2793       LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN        BUILTIN        FPC CPU
  PIC 0                  BUILTIN        BUILTIN        12x100GE/36x40GE/144x10GE

    Xcvr 0       REV 01   740-032986     QH0400VH       QSFP+-40G-SR4
    Xcvr 1       REV 01   740-032986     QH0400VM       QSFP+-40G-SR4
    Xcvr 35      REV 01   740-058734     1ECQ11390ZB    QSFP-100GBASE-SR4
FPC 8            REV 36   750-077140     ACNP4625       LC1102 - 12C / 36Q /
144X
  CPU                     BUILTIN        BUILTIN        FPC CPU
  PIC 0                  BUILTIN        BUILTIN        12x100GE/36x40GE/144x10GE

    Xcvr 1       REV 01   740-058732     1AMQA14206D    QSFP-100GBASE-LR4
    Xcvr 10      REV 01   740-032986     QF4301KB       QSFP+-40G-SR4
    Xcvr 24      REV 01   740-054050     INFAJ0492244   QSFP+-4X10G-LR
FPC 9           REV 35   750-071976     ACPD3055       LC1101 - 30C / 30Q / 96X

  CPU                     BUILTIN        BUILTIN        FPC CPU
  PIC 0                   BUILTIN        BUILTIN        30x100GE/30x40GE/96x10GE

    Xcvr 0                NON-JNPR       INGBT7970007   QSFP-100GBASE-LR4
    Xcvr 1                NON-JNPR       UWQ24D9        QSFP-100GBASE-LR4
    Xcvr 2                NON-JNPR       INGBT7970011   QSFP-100GBASE-LR4
    Xcvr 3                NON-JNPR       UX60AF1        QSFP-100G-CWDM4
    Xcvr 4                NON-JNPR       UX408JJ        QSFP-100GBASE-LR4
```



```
    Xcvr 11        REV 01    740-058734    1ECQ113835F       QSFP-100GBASE-SR4
    Xcvr 18                  NON-JNPR      Q7496             QSFP-100G-CWDM4
    Xcvr 29        REV 01    740-058734    1ECQ11380LZ       QSFP-100GBASE-SR4
Power Supply 0     REV 02    740-049388    1EDL625039E       Power Supply AC
Power Supply 1     REV 02    740-049388    1EDL62503AD       Power Supply AC
Power Supply 2     REV 02    740-049388    1EDL625039P       Power Supply AC
Power Supply 3     REV 02    740-049388    1EDL702004E       Power Supply AC
Power Supply 4     REV 02    740-049388    1EDL625039D       Power Supply AC
Power Supply 5     REV 02    740-049388    1EDL63706JD       Power Supply AC
Power Supply 6     REV 02    740-049388    1EDL63706JH       Power Supply AC
FTC 0              REV 10    750-050309    ACPM2918          Fan Controller 16
FTC 1              REV 10    750-050309    ACPE8185          Fan Controller 16
Fan Tray 0         REV 10    760-077141    ACPV7288          Fan Tray 16
Fan Tray 1         REV 10    760-057901    ACPL0546          Fan Tray 16
SIB 0              REV 15    750-058270    ACPM2804          Switch Fabric 16
SIB 1              REV 15    750-058270    ACPM2808          Switch Fabric 16
SIB 2              REV 15    750-058270    ACPL4450          Switch Fabric 16
SIB 3              REV 15    750-058270    ACPJ9834          Switch Fabric 16
SIB 4              REV 15    750-058270    ACPM2814          Switch Fabric 16
SIB 5              REV 15    750-058270    ACPL4277          Switch Fabric 16
FPD Board          REV 07    711-054687    ACPL1407          Front Panel Display
```

show chassis hardware (M7i Router)

```
user@host> show chassis hardware
Hardware inventory:
Item               Version   Part number   Serial number     Description
Chassis                                    31959             M7i
Midplane           REV 02    710-008761    CA0209            M7i Midplane
Power Supply 0     Rev 04    740-008537    PD10272           AC Power Supply
Routing Engine     REV 01    740-008846    1000396803        RE-5.0
CFEB               REV 02    750-009492    CA0166            Internet Processor IIv1
FPC 0                                                        E-FPC
  PIC 0            REV 04    750-003163    HJ6416            1x G/E, 1000 BASE-SX
  PIC 1            REV 04    750-003163    HJ6423            1x G/E, 1000 BASE-SX
  PIC 2            REV 04    750-003163    HJ6421            1x G/E, 1000 BASE-SX
  PIC 3            REV 02    750-003163    HJ0425            1x G/E, 1000 BASE-SX
FPC 1                                                        E-FPC
  PIC 2            REV 01    750-009487    HM2275            ASP - Integrated
  PIC 3            REV 01    750-009098    CA0142            2x F/E, 100 BASE-TX

Hardware inventory:
Item               Version   Part number   Serial number     Description
Chassis                                    B1157             M7i
Midplane           REV 05    710-008761    DM0840            M7i Midplane
Power Supply 0     Rev 08    740-008537    TE53755           AC Power Supply
Routing Engine     REV 07    740-011202    1000736567        RE-850
CFEB               REV 09    750-010463    DK6952            Internet Processor II
FPC 0                                                        E-FPC
  PIC 0            REV 12    750-012838    DL7993            4x 1GE(LAN), IQ2
    Xcvr 0         REV 01    740-011614    PD94TDJ           SFP-LX10
    Xcvr 1         REV 01    740-011615    PAD5EER           UNSUPPORTED
    Xcvr 2         REV 01    740-011614    PD94THU           SFP-LX10
    Xcvr 3                   NON-JNPR      PDC2E7A           SFP-LX10
  PIC 1            REV 03    750-023116    JT0203            4x CHSTM1 SDH CE SFP
    Xcvr 0         REV 01    740-012434    AGT063832PS       SFP-SR
    Xcvr 1         REV 01    740-012434    AGT063832LY       SFP-SR
    Xcvr 3         REV 01    740-016064    C06J19018         SFP-LR
  PIC 2            REV 15    750-014895    DM5757            MultiServices 100
  PIC 3            REV 01    750-025390    JW9448            12x T1/E1 CE
FPC 1                                                        E-FPC
```

```
  PIC 2                   BUILTIN      BUILTIN           1x Tunnel
  PIC 3          REV 09   750-009099   DM0899            1x G/E, 1000 BASE
    Xcvr 0       REV 01   740-012434   AGT07150HGJ       UNSUPPORTED
Fan Tray                                                 Rear Fan Tray
```

show chassis hardware (M10 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                1122              M10
Midplane         REV 1.1  710-001950   S/N AC6626
Power supply A   Rev 01   740-002497   S/N LC36095       AC
Power supply B   Rev 01   740-002497   S/N LC36100       AC
Display          REV 1.2  710-001995   S/N AC6656
Host                                   18000005dfb3fb01  teknor
FEB              REV 01   710-001948   S/N AC6632        Internet Processor II
FPC 0
  PIC 0          REV 08   750-001072   S/N AB2485        1x G/E, 1000 BASE-SX
  PIC 1          REV 01   750-000613   S/N AA1048        1x OC-12 SONET, SMIR
FPC 1
Fan Tray 0                                               FANTRAY-M10I-S
Fan Tray 1                                               FANTRAY-M10I-S
```

show chassis hardware models (M10 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 04   710-008920                     CHAS-MP-M10i-S
Power Supply 0   Rev 06   740-008537                     PWR-M10i-M7i-AC-S
Power Supply 1   Rev 06   740-008537                     PWR-M10i-M7i-AC-S
HCM 0            REV 03   710-010580                     HCM-M10i-S
HCM 1            REV 03   710-010580                     HCM-M10i-S
Routing Engine 0 REV 09   740-009459                     RE-400-256-S
CFEB 0           REV 05   750-010465                     FEB-M10i-M7i-S
FPC 0
  PIC 0          REV 10   750-002971                     PE-4OC3-SON-MM
  PIC 1          REV 11   750-002992                     PE-4FE-TX
  PIC 2          REV 03   750-002977                     PE-2OC3-ATM-MM
  PIC 3          REV 08   750-005724                     PE-2OC3-ATM2-MM
FPC 1
  PIC 2          REV 12   750-008425                     PE-AS
  PIC 3          REV 13   750-005636                     PE-4CHDS3-QPP
Fan Tray 0                                               FANTRAY-M10I-S
Fan Tray 1                                               FANTRAY-M10I-S
```

show chassis hardware (M20 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                20033             M20
Backplane        REV 07   710-001517   S/N AA7940
Power supply B   Rev 01   740-001465   S/N 000001        AC
Display          REV 02   710-001519   S/N AA9704
Host 0                                 98000004f8f27501  teknor
SSB slot 0       REV 01   710-001951   S/N AD5905        Internet Processor II
  SSRAM bank 0   REV 01   710-001385   S00480            2 MB
  SSRAM bank 1   REV 01   710-001385   S00490            2 MB
```



```
  SSRAM bank 2    REV 01   710-001385   S001:?            2 MB
  SSRAM bank 3    REV 01   710-001385   S00483            2 MB
SSB slot 1        N/A      N/A          N/A               Backup
FPC 1             REV 01   710-001292   S/N AB7528
  SSRAM           REV 01   710-000077   S/N 304209        1 MB
  SDRAM bank 0    REV 01   710-000099   S/N 000603        64 MB
  SDRAM bank 1    REV 01   710-000099   S/N 000414        64 MB
  PIC 0           REV 03   750-000612   S/N AB8433        2x OC-3 ATM, MM
  PIC 1           REV 01   750-000616   S/N AA1168        1x OC-12 ATM, MM
  PIC 2           REV 01   750-000613   S/N AA1008        1x OC-12 SONET, SMIR
  PIC 3           REV 01   750-002501   S/N AD5810        4x E3
FPC 2             REV 01   710-001292   S/N AC0119
  SSRAM           REV 01   710-000077   S/N 503241        1 MB
  SDRAM bank 0    REV 01   710-000099   S/N 306835        64 MB
  SDRAM bank 1    REV 01   710-000099   S/N 306832        64 MB
Fan Tray 0                                                Front Upper Fan Tray
Fan Tray 1                                                Front Middle Fan Tray
Fan Tray 2                                                Front Bottom Fan Tray
Fan Tray 3                                                Rear Fan Tray
```

show chassis hardware models (M20 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item              Version  Part number  CLEI code         FRU model number
Backplane         REV 03   710-002334                     CHAS-MP-M20-S
Power Supply A    REV 06   740-001465                     PWR-M20-AC-S
Display           REV 04   710-001519                     CRAFT-M20-S
Routing Engine 0  REV 06   740-003239                     RE-333-768-S
Routing Engine 1  REV 06   740-003239                     RE-333-768-S
SSB 0             REV 02   710-001951                     SSB-E-M20
SSB 1             N/A      N/A
FPC 0             REV 03   710-003308                     FPC-E
  PIC 0           REV 08   750-002303                     P-4FE-TX
  PIC 1           REV 07   750-004745                     P-2MCDS3
  PIC 2           REV 03   750-002965                     PE-4CHDS3
FPC 1             REV 03   710-003308                     FPC-E
  PIC 0           REV 03   750-002914                     P-2OC3-ATM-MM
Fan Tray 0                                                FANTRAY-F-M20-S
Fan Tray 1                                                FANTRAY-F-M20-S
Fan Tray 2                                                FANTRAY-F-M20-S
Fan Tray 3                                                FANTRAY-R-M20-S
```

show chassis hardware (M40 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number     Description
Backplane         REV 02   710-000073   S/N AA0053
Power supply A    Rev 2    740-000235   S/N 000042        DC
Maxicab           REV X1   710-000229   S/N AA0139
Minicab           REV X1   710-000482   S/N AA0201
Display           REV 06   710-000150   S/N AA0905
Host                                                      cpv5000
SCB               REV X1   710-000075   S/N AA0158        Internet Processor I
  SSRAM bank 0    REV 02   710-000077   S/N AA2267        1 MB
  SSRAM bank 1    REV 02   710-000077   S/N AA2270        1 MB
  SSRAM bank 2    REV 02   710-000077   S/N AA2269        1 MB
  SSRAM bank 3    REV 02   710-000077   S/N AA2268        1 MB
FPC 0             REV 01   710-000175   S/N AA0048
```



```
  SSRAM          REV 01   710-000077   S/N AA2333        1 MB
  SDRAM bank 0   REV 01   710-000099   S/N AA2332        64 MB
  SDRAM bank 1   REV X1   710-000099   S/N AA2337        64 MB
  PIC 0          REV 04   750-000613   S/N aa0343        1x OC-12 SONET, SMIR
  PIC 1          REV 04   750-000613   S/N AA0379        1x OC-12 SONET, SMIR
  PIC 2          REV 04   750-000613   S/N AA0377        1x OC-12 SONET, SMIR
  PIC 3          REV 04   750-000613   S/N AA0378        1x Tunnel
FPC 2            REV 01   710-000175   S/N AA0042
  SSRAM          REV 02   710-000077   S/N AA2288        1 MB
  SDRAM bank 0   REV 01   710-000099   S/N AA2331        64 MB
  SDRAM bank 1   REV 01   710-000099   S/N AA2330        64 MB
  PIC 0          REV X1   750-000603   S/N AA0143        4x OC-3 SONET, SMIR
  PIC 1          REV X1   750-000615   S/N AA0149        4x OC-3 SONET, MM
  PIC 2          REV X1   750-000611   S/N AA0148        4x OC-3 SONET, MM
  PIC 3          REV 04   750-000613   S/N AA0330        1x OC-12 SONET, SMIR
FPC 4            REV 01   710-000175   S/N AA0050
  SSRAM          REV 01   710-000077   S/N AA2327        1 MB
  SDRAM bank 0   REV 01   710-000099   S/N AA2329        64 MB
  SDRAM bank 1   REV 01   710-000099   S/N AA2328        64 MB
  PIC 0          REV 04   750-000613   S/N AA0320        1x OC-12 SONET, SMIR
  PIC 2          REV 05   750-000616   S/N AA1341        1x OC-12 ATM, MM
  PIC 3          REV 08   750-001072   S/N AB2462        1x G/E, 1000 BASE-SX
FPC 5            REV 10   710-000175   S/N AA7663
  SSRAM          REV 01   710-000077   S/N 501590        1 MB
  SDRAM bank 0   REV 01   710-000099   S/N 300949        64 MB
  SDRAM bank 1   REV 01   710-000099   S/N 300868        64 MB
  PIC 1          REV 01   750-001323   S/N AB1670        1x Tunnel
```

show chassis hardware (M40e Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                                  m40e
Midplane         REV 01   710-005071   AX3671
FPM CMB          REV 03   710-001642   AR9074
FPM Display      REV 03   710-001647   AR7331
CIP              REV 04   710-002649   BB4449
PEM 0            Rev 01   740-003787   MC12364           Power Entry Module
PEM 1            Rev 01   740-003787   MC12383           Power Entry Module
PCG 0            REV 07   710-001568   AG1332
PCG 1            REV 07   710-001568   AR3789
Host 0                                 3e000007c8176601  Present
MCS 0            REV 11   710-001226   AN5813
SFM 0 SPP        REV 07   710-001228   AG4676
SFM 0 SPR        REV 05   710-002189   AE4735            Internet Processor II
SFM 1 SPP        REV 07   710-001228   AP1347
SFM 1 SPR        REV 05   710-002189   BE0063            Internet Processor II
FPC 0            REV 01   710-011725   BE0669            M40e-EP-FPC Type 1
  CPU            REV 01   710-004600   BD9504
  PIC 0          REV 03   750-003737   AY3991            4x G/E, 1000 BASE-SX
FPC 1            REV 01   710-005197   BD9842            M40e-FPC Type 2
  CPU            REV 01   710-004600   BB4869
  PIC 0          REV 07   750-001900   AR8278            1x OC-48 SONET, SMSR
FPC 2            REV 02   710-005197   BD9824            M40e-FPC Type 2
  CPU            REV 01   710-004600   BD9531
  PIC 0          REV 03   750-003737   AY3986            4x G/E, 1000 BASE-SX
FPC 4            REV 02   710-005078   BE0664            M40e-FPC Type 1
  CPU            REV 01   710-004600   BD9559
  PIC 0          REV 03   750-001894   AG7963            1x G/E, 1000 BASE-SX
  PIC 2          REV 01   750-002575   AF2472            4x OC-3 SONET, SMIR
```

```
FPC 6            REV 02   710-005078   BE0652           M40e-FPC Type 1
  CPU            REV 01   710-004600   BD9607
  PIC 0          REV 02   750-002911   AN2286           4x F/E, 100 BASE-TX
  PIC 2          REV 01   750-002577   AP6345           4x OC-3 SONET, MM
```

show chassis hardware (M120 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN000054AC       M120
Midplane         REV 01   710-013667   RB4170           M120 Midplane
FPM Board        REV 02   710-011407   CJ9186           M120 FPM Board
FPM Display      REV 02   710-011405   CJ9173           M120 FPM Display
FPM CIP          REV 02   710-011410   CJ9221           M120 FPM CIP
PEM 0            Rev 05   740-011936   RM28320          AC Power Entry Module
PEM 1            Rev 05   740-011936   RM28321          AC Power Entry Module
Routing Engine 0 REV 03   740-014080   1000642883       RE-A-1000
CB 0             REV 03   710-011403   CM8346           M120 Control Board
CB 1             REV 06   710-011403   CP6728           M120 Control Board
FPC 1            REV 02   710-015908   CP6925           M120 CFPC 10GE
  PIC 0                   BUILTIN      BUILTIN          1x 10GE(LAN/WAN) XFP
    Xcvr 0       REV 01   740-014279   62E204N00007     XFP-10G-LR
FPC 3            REV 03   710-011393   CJ9234           M120 FPC Type 2
  PIC 0          REV 16   750-008155   NB5229           2x G/E IQ, 1000 BASE
    Xcvr 0       REV 01   740-011613   P9F15JB          SFP-SX
    Xcvr 1       REV 01   740-007326   P4Q0R9G          SFP-SX
  PIC 1          REV 09   750-007745   CG4360           4x OC-3 SONET, SMIR
  PIC 2          REV 16   750-008155   ND7787           2x G/E IQ, 1000 BASE
    Xcvr 0       REV 01   740-011613   P9F12AS          SFP-SX
    Xcvr 1       REV 01   740-011613   P9F1ALU          SFP-SX
  PIC 3          REV 07   750-011800   JW1284           8x 1GE(LAN), IQ2
    Xcvr 0       REV 01   740-011613   P9F1AM6          SFP-SX
    Xcvr 6       REV 01   740-011613   P9F16NN          SFP-SX
    Xcvr 7       REV 01   740-011782   P8C29Y7          SFP-SX
  Board B        REV 02   710-011395   CN3754           M120 FPC Mezz
FPC 4            REV 02   710-011398   CP6741           M120 FPC Type 3
  PIC 0          REV 16   750-007141   NB2855           10x 1GE(LAN), 1000 BASE

    Xcvr 0       REV 01   740-011782   P922A1F          SFP-SX
    Xcvr 1       REV 01   740-011782   P922A16          SFP-SX
    Xcvr 2       REV 01   740-011782   P922A0U          SFP-SX
    Xcvr 3       REV 01   740-011782   P9229UZ          SFP-SX
    Xcvr 4       REV 01   740-009029   P11JXWP          SFP-LX
    Xcvr 6       REV 01   740-011613   P9F1ALW          SFP-SX
FPC 5            REV 01   710-011388   CJ9088           M120 FPC Type 1
  PIC 0          *** Hardware Not Supported ***
  PIC 1          REV 05   750-012052   NB0410           1x CHOC3 IQ SONET, SMLR

  PIC 2          REV 01   750-013167   CM3824           4x CHDS3 IQ
  PIC 3          REV 01   750-010240   CB5366           1x G/E SFP, 1000 BASE
  Board B        REV 01   710-011390   CJ9103           M120 FPC Mezz Board
FEB 3            REV 04   710-011663   CP6673           M120 FEB
FEB 4            REV 04   710-011663   CJ9368           M120 FEB
FEB 5            REV 04   710-011663   CJ9386           M120 FEB
Fan Tray 0                                              Front Top Fan Tray
Fan Tray 1                                              Front Bottom Fan Tray
Fan Tray 2                                              Rear Top Fan Tray
Fan Tray 3                                              Rear Bottom Fan Tray
```


show chassis hardware detail (M120 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN000054AC        M120
Midplane         REV 01   710-013667   RB4170            M120 Midplane
FPM Board        REV 02   710-011407   CJ9186            M120 FPM Board
FPM Display      REV 02   710-011405   CJ9173            M120 FPM Display
FPM CIP          REV 02   710-011410   CJ9221            M120 FPM CIP
PEM 0            Rev 05   740-011936   RM28320           AC Power Entry Module
PEM 1            Rev 05   740-011936   RM28321           AC Power Entry Module
Routing Engine 0 REV 03   740-014080   1000642883        RE-A-1000
  ad0      248 MB  SILICONSYSTEMS INC 256M 126CT505S0763SC00110 Compact Flash
  ad2    38154 MB  HTE541040G9SA00     MPBBT0X2HS2E3M    Hard Disk
CB 0             REV 03   710-011403   CM8346            M120 Control Board
CB 1             REV 06   710-011403   CP6728            M120 Control Board
FPC 1            REV 02   710-015908   CP6925            M120 CFPC 10GE
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) XFP
    Xcvr 0       REV 01   740-014279   62E204N00007      XFP-10G-LR
FPC 3            REV 03   710-011393   CJ9234            M120 FPC Type 2
  PIC 0          REV 16   750-008155   NB5229            2x G/E IQ, 1000 BASE
    Xcvr 0       REV 01   740-011613   P9F15JB           SFP-SX
    Xcvr 1       REV 01   740-007326   P4Q0R9G           SFP-SX
  PIC 1          REV 09   750-007745   CG4360            4x OC-3 SONET, SMIR
  PIC 2          REV 16   750-008155   ND7787            2x G/E IQ, 1000 BASE
    Xcvr 0       REV 01   740-011613   P9F12AS           SFP-SX
    Xcvr 1       REV 01   740-011613   P9F1ALU           SFP-SX
  PIC 3          REV 07   750-011800   JW1284            8x 1GE(LAN), IQ2
    Xcvr 0       REV 01   740-011613   P9F1AM6           SFP-SX
    Xcvr 6       REV 01   740-011613   P9F16NN           SFP-SX
    Xcvr 7       REV 01   740-011782   P8C29Y7           SFP-SX
  Board B        REV 02   710-011395   CN3754            M120 FPC Mezz
FPC 4            REV 02   710-011398   CP6741            M120 FPC Type 3
  PIC 0          REV 16   750-007141   NB2855            10x 1GE(LAN), 1000 BASE

    Xcvr 0       REV 01   740-011782   P922A1F           SFP-SX
    Xcvr 1       REV 01   740-011782   P922A16           SFP-SX
    Xcvr 2       REV 01   740-011782   P922A0U           SFP-SX
    Xcvr 3       REV 01   740-011782   P9229UZ           SFP-SX
    Xcvr 4       REV 01   740-009029   P11JXWP           SFP-LX
    Xcvr 6       REV 01   740-011613   P9F1ALW           SFP-SX
FPC 5            REV 01   710-011388   CJ9088            M120 FPC Type 1
  PIC 0          *** Hardware Not Supported ***
  PIC 1          REV 05   750-012052   NB0410            1x CHOC3 IQ SONET, SMLR

  PIC 2          REV 01   750-013167   CM3824            4x CHDS3 IQ
  PIC 3          REV 01   750-010240   CB5366            1x G/E SFP, 1000 BASE
  Board B        REV 01   710-011390   CJ9103            M120 FPC Mezz Board
FEB 3            REV 04   710-011663   CP6673            M120 FEB
FEB 4            REV 04   710-011663   CJ9368            M120 FEB
FEB 5            REV 04   710-011663   CJ9386            M120 FEB
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Top Fan Tray
Fan Tray 3                                               Rear Bottom Fan Tray
```

show chassis hardware models (M120 Router)

```
user@host> show chassis hardware models
```



```
Hardware inventory:
Item             Version  Part number  CLEI code          FRU model number
Midplane         REV 01   710-013667
FPM CIP          REV 02   710-011410                      CRAFT-M120-S
PEM 0            Rev 05   740-011936                      PWR-M120-AC-S
PEM 1            Rev 05   740-011936                      PWR-M120-AC-S
Routing Engine 0 REV 03   740-014080                      RE-A-1000-2048-S
CB 0             REV 03   710-011403                      CB-M120-S
CB 1             REV 06   710-011403                      CB-M120-S
FPC 1            REV 02   710-015908                      M120-cFPC-1XGE-XFP
FPC 3
  PIC 0          REV 16   750-008155                      PB-2GE-SFP-QPP
  PIC 1          REV 09   750-007745                      PC-4OC3-SON-SMIR
  PIC 2          REV 16   750-008155                      PB-2GE-SFP-QPP
  PIC 3          REV 07   750-011800                      PB-8GE-TYPE2-SFP-IQ2
FPC 4
  PIC 0          REV 16   750-007141                      PC-10GE-SFP
FPC 5
  PIC 1          REV 05   750-012052                      PB-1CHOC3-SMIR-QPP
  PIC 2          REV 01   750-013167                      PE-4CHDS3-QPP
  PIC 3          REV 01   750-010240                      PB-1GE-SFP
Fan Tray 0                                                FFANTRAY-M120-S
Fan Tray 1                                                FFANTRAY-M120-S
Fan Tray 2                                                RFANTRAY-M120-S
Fan Tray 3                                                RFANTRAY-M120-S
```

show chassis hardware (M160 Router)

```
user@host> show chassis hardware
Item             Version  Part number  Serial number      Description
Chassis                                101                M160
Midplane         REV 02   710-001245   S/N AB4107
FPM CMB          REV 01   710-001642   S/N AA2911
FPM Display      REV 01   710-001647   S/N AA2999
CIP              REV 02   710-001593   S/N AA9563
PEM 0            Rev 01   740-001243   S/N KJ35769        DC
PEM 1            Rev 01   740-001243   S/N KJ35765        DC
PCG 0            REV 01   710-001568   S/N AA9794
PCG 1            REV 01   710-001568   S/N AA9804
Host 1                                 da000004f8d57001   teknor
MCS 1            REV 03   710-001226   S/N AA9777
SFM 0 SPP        REV 04   710-001228   S/N AA2975
SFM 0 SPR        REV 02   710-001224   S/N AA9838         Internet Processor I
SFM 1 SPP        REV 04   710-001228   S/N AA2860
SFM 1 SPR        REV 01   710-001224   S/N AB0139         Internet Processor I
FPC 0            REV 03   710-001255   S/N AA9806         FPC Type 1
  CPU            REV 02   710-001217   S/N AA9590
  PIC 1          REV 05   750-000616   S/N AA1527         1x OC-12 ATM, MM
  PIC 2          REV 05   750-000616   S/N AA1535         1x OC-12 ATM, MM
  PIC 3          REV 01   750-000616   S/N AA1519         1x OC-12 ATM, MM
FPC 1            REV 02   710-001611   S/N AA9523         FPC Type 2
  CPU            REV 02   710-001217   S/N AA9571
  PIC 0          REV 03   750-001900   S/N AA9626         1x STM-16 SDH, SMIR
  PIC 1          REV 01   710-002381   S/N AD3633         2x G/E, 1000 BASE-SX
FPC 2                                                     FPC Type OC192
  CPU            REV 03   710-001217   S/N AB3329
  PIC 0          REV 01                                   1x OC-192 SM SR-2
Fan Tray 0                                                Rear Bottom Blower
Fan Tray 1                                                Rear Top Blower
Fan Tray 2                                                Front Top Blower
Fan Tray 3                                                Front Fan Tray
```


show chassis hardware models (M160 Router)

```
user@host> show chassis hardware models
 Hardware inventory:
Item             Version  Part number  CLEI code        FRU model number
Midplane         REV 03   710-009120                    CHAS-BP-M320-S
FPM Display      REV 02   710-009351                    CRAFT-M320-S
CIP              REV 03   710-005926                    CIP-M320-S
PEM 2            Rev X4   740-009148                    PWR-M-DC-S
PEM 3            Rev X4   740-009148                    PWR-M-DC-S
Routing Engine 0 REV 02   740-008883                    RE-1600-2048-S
Routing Engine 1 REV 02   740-008883                    RE-1600-2048-S
FPC 0            REV 02   710-010419                    M320-FPC1
  PIC 0          REV 01   750-001323                    P-TUNNEL
  PIC 1          REV 02   750-002987                    PE-1OC12-SON-SMIR
  PIC 2          REV 04   750-001894                    PB-1GE-SX
  PIC 3          REV 04   750-001896                    PB-1OC12-SON-SMIR
FPC 1            REV 02   710-010419                    M320-FPC1
  PIC 0          REV 04   750-001894                    PB-1GE-SX
  PIC 1          REV 04   750-001894                    PB-1GE-SX
  PIC 3          REV 03   750-001894                    PB-1GE-SX
FPC 2            REV 02   710-010419                    M320-FPC1
  PIC 0          REV 10   750-005634                    PB-1CHOC12SMIR-QPP
  PIC 1          REV 10   750-005634                    PB-1CHOC12SMIR-QPP
  PIC 2          REV 07   750-005634                    PB-1CHOC12SMIR-QPP
  PIC 3          REV 07   750-005634                    PB-1CHOC12SMIR-QPP
  PIC 1          REV 10   750-005634                    PB-1CHOC12SMIR-QPP
  PIC 2          REV 07   750-005634                    PB-1CHOC12SMIR-QPP
  PIC 3          REV 07   750-005634                    PB-1CHOC12SMIR-QPP
FPC 3
  PIC 0          REV 03   750-001895                    PB-1OC12-SON-MM
  PIC 1          REV 04   750-001894                    PB-1GE-SX
  PIC 3          REV 04   750-003141                    PB-1GE-SX-B
FPC 4            REV 02   710-010419                    M320-FPC1
FPC 5            REV 02   710-010419                    M320-FPC1
FPC 6            REV 02   710-010419                    M320-FPC1
FPC 7
  PIC 0          REV 15   750-001901                    PB-4OC12-SON-SMIR
  PIC 1          REV 06   750-001900                    PB-1OC48-SON-SMSR
  PIC 2          REV 07   750-001900                    PB-1OC48-SON-SMSR
  PIC 3          REV 05   750-003737                    PB-4GE-SX
SIB 0            REV 03   710-009184                    SIB-M-S
SIB 1            REV 03   710-009184                    SIB-M-S
SIB 2            REV 03   710-009184                    SIB-M-S
SIB 3            REV 03   710-009184                    SIB-M-S
Fan Tray 0                                              FFANTRAY-M320-S
Fan Tray 1                                              FFANTRAY-M320-S
Fan Tray 2                                              RFANTRAY-M320-S
```

show chassis hardware detail (M160 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                101              M160
Midplane         REV 02   710-001245   S/N AB4107
FPM CMB          REV 01   710-001642   S/N AA2911
FPM Display      REV 01   710-001647   S/N AA2999
CIP              REV 02   710-001593   S/N AA9563
PEM 0            Rev 01   740-001243   S/N KJ35769      DC
```

```
PEM 1            Rev 01   740-001243   S/N KJ35765       DC
PCG 0            REV 01   710-001568   S/N AA9794
PCG 1            REV 01   710-001568   S/N AA9804
Host 1                                 da000004f8d57001  teknor
MCS 1            REV 03   710-001226   S/N AA9777
SFM 0 SPP        REV 04   710-001228   S/N AA2975
SFM 0 SPR        REV 02   710-001224   S/N AA9838        Internet Processor I
  SSRAM bank 0   REV 01   710-000077   S/N 306456        1 MB
  SSRAM bank 1   REV 01   710-000077   S/N 306474        1 MB
  SSRAM bank 2   REV 01   710-000077   S/N 306388        1 MB
  SSRAM bank 3   REV 01   710-000077   S/N 306392        1 MB
SFM 1 SPP        REV 04   710-001228   S/N AA2860
SFM 1 SPR        REV 01   710-001224   S/N AB0139        Internet Processor I
  SSRAM bank 0   REV 01   710-000077   S/N 302917        1 MB
  SSRAM bank 1   REV 01   710-000077   S/N 302662        1 MB
  SSRAM bank 2   REV 01   710-000077   S/N 302593        1 MB
  SSRAM bank 3   REV 01   710-000077   S/N 100160        1 MB
FPC 0            REV 03   710-001255   S/N AA9806        FPC Type 1
  CPU            REV 02   710-001217   S/N AA9590
  SSRAM          REV 01   710-000077   S/N 302836        1 MB
  SDRAM 0        REV 01   710-001196   S00141            32 MB
  SDRAM 1        REV 01   710-001196   S0010;            32 MB
  SSRAM          REV 01   710-000077   S/N 302633        1 MB
  SDRAM 0        REV 01   710-001196   S00143            32 MB
  SDRAM 1        REV 01   710-001196   S00115            32 MB
  SSRAM          REV 01   710-000077   S/N 302952        1 MB
  SDRAM 0        REV 01   710-001196   S00135            32 MB
  SDRAM 1        REV 01   710-001196   S001=3            32 MB
  SSRAM          REV 01   710-000077   S/N 302892        1 MB
  SDRAM 0        REV 01   710-001196   S000?6            32 MB
  SDRAM 1        REV 01   710-001196   S001=5            32 MB
  PIC 1          REV 05   750-000616   S/N AA1527        1x OC-12 ATM, MM
  PIC 2          REV 05   750-000616   S/N AA1535        1x OC-12 ATM, MM
  PIC 3          REV 01   750-000616   S/N AA1519        1x OC-12 ATM, MM
FPC 1            REV 02   710-001611   S/N AA9523        FPC Type 2
  CPU            REV 02   710-001217   S/N AA9571
  SSRAM          REV 01   710-000077   S/N 306340        1 MB
  SDRAM 0        REV 01   710-001196   S00012            32 MB
  SDRAM 1        REV 01   710-001196   S0001?            32 MB
  SSRAM          REV 01   710-000077   S/N 306454        1 MB
  SDRAM 0        REV 01   710-001196   S00028            32 MB
  SDRAM 1        REV 01   710-001196   S0002?            32 MB
  SSRAM          REV 01   710-000077   S/N 306492        1 MB
  SDRAM 0        REV 01   710-001196   S00015            32 MB
  SDRAM 1        REV 01   710-001196   S00031            32 MB
  SSRAM          REV 01   710-000077   S/N 306363        1 MB
  SDRAM 0        REV 01   710-001196   S00013            32 MB
  SDRAM 1        REV 01   710-001196   S00032            32 MB
  PIC 0          REV 03   750-001900   S/N AA9626        1x STM-16 SDH, SMIR
  PIC 1          REV 01   710-002381   S/N AD3633        2x G/E, 1000 BASE-SX
FPC 2                                                    FPC Type OC192
... SSRAM          REV 01   710-000077   S/N 306466        1 MB
```

show chassis hardware (M320 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                67245             M320
Midplane         REV 05   710-009120   RB1202            M320 Midplane
FPM GBUS         REV 04   710-005928   HZ5697            M320 Board
```



```
FPM Display       REV 05   710-009351   HR1464          M320 FPM Display
CIP               REV 04   710-005926   HT8672          M320 CIP
PEM 0             Rev 05   740-009148   QK34208         DC Power Entry Module
PEM 1             Rev 05   740-009148   QK34262         DC Power Entry Module
PEM 2             Rev 05   740-009148   QF10449         DC Power Entry Module
PEM 3             Rev 05   740-009148   QJ18257         DC Power Entry Module
Routing Engine 0 REV 06   740-008883   P11123901185    RE-4.0
CB 0              REV 07   710-009115   JB2382          M320 Control Board
FPC 0             REV 02   710-005017   CD9926          M320 FPC Type 2
  CPU             REV 01   710-011659   CJ6940          M320 PCA SCPU
  PIC 0           REV 07   750-001900   AT1594          1x OC-48 SONET, SMSR
  PIC 1           REV 03   750-001850   HS2746          1x Tunnel
  PIC 2           REV 05   750-010618   JE7117          4x G/E SFP, 1000 BASE
  PIC 3           REV 06   750-001900   HE6083          1x OC-48 SONET, SMSR
FPC 2             REV 02   710-005017   CH0319          M320 FPC Type 1
  CPU             REV 01   710-011659   CJ6942          M320 PCA SCPU
  PIC 0           REV 05   750-003034   BD8705          4x OC-3 SONET, SMIR
FPC 5             REV 02   710-005017   CD9938          M320 FPC Type 2
  CPU
FPC 7             REV 02   710-005017   CD9934          M320 FPC Type 2
  CPU
SIB 0             REV 09   710-009184   JA6540          M320 SIB
SIB 1             REV 09   710-009184   HV9511          M320 SIB
SIB 2             REV 09   710-009184   HW2057          M320 SIB
SIB 3             REV 09   710-009184   JA6687          M320 SIB
Fan Tray 0                                              Front Top Fan Tray
Fan Tray 1                                              Front Bottom Fan Tray
Fan Tray 2                                              Rear Fan Tray
```

show chassis hardware models (M320 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item              Version  Part number  CLEI code       FRU model number
Midplane          REV 03   710-009120                   CHAS-BP-M320-S
FPM Display       REV 02   710-009351                   CRAFT-M320-S
CIP               REV 03   710-005926                   CIP-M320-S
PEM 2             Rev X4   740-009148                   PWR-M-DC-S
PEM 3             Rev X4   740-009148                   PWR-M-DC-S
Routing Engine 0 REV 02   740-008883                   RE-1600-2048-S
Routing Engine 1 REV 02   740-008883                   RE-1600-2048-S
FPC 0             REV 02   710-010419                   M320-FPC1
  PIC 0           REV 01   750-001323                   P-TUNNEL
  PIC 1           REV 02   750-002987                   PE-1OC12-SON-SMIR
  PIC 2           REV 04   750-001894                   PB-1GE-SX
  PIC 3           REV 04   750-001896                   PB-1OC12-SON-SMIR
FPC 1             REV 02   710-010419                   M320-FPC1
  PIC 0           REV 04   750-001894                   PB-1GE-SX
  PIC 1           REV 04   750-001894                   PB-1GE-SX
  PIC 3           REV 03   750-001894                   PB-1GE-SX
FPC 2             REV 02   710-010419                   M320-FPC1
  PIC 0           REV 10   750-005634                   PB-1CHOC12SMIR-QPP
  PIC 1           REV 10   750-005634                   PB-1CHOC12SMIR-QPP
  PIC 2           REV 07   750-005634                   PB-1CHOC12SMIR-QPP
  PIC 3           REV 07   750-005634                   PB-1CHOC12SMIR-QPP
  PIC 1           REV 10   750-005634                   PB-1CHOC12SMIR-QPP
  PIC 2           REV 07   750-005634                   PB-1CHOC12SMIR-QPP
  PIC 3           REV 07   750-005634                   PB-1CHOC12SMIR-QPP
FPC 3
  PIC 0           REV 03   750-001895                   PB-1OC12-SON-MM
  PIC 1           REV 04   750-001894                   PB-1GE-SX
```



```
  PIC 3          REV 04   750-003141                      PB-1GE-SX-B
FPC 4            REV 02   710-010419                      M320-FPC1
FPC 5            REV 02   710-010419                      M320-FPC1
FPC 6            REV 02   710-010419                      M320-FPC1
FPC 7
  PIC 0          REV 15   750-001901                      PB-4OC12-SON-SMIR
  PIC 1          REV 06   750-001900                      PB-1OC48-SON-SMSR
  PIC 2          REV 07   750-001900                      PB-1OC48-SON-SMSR
  PIC 3          REV 05   750-003737                      PB-4GE-SX
SIB 0            REV 03   710-009184                      SIB-M-S
SIB 1            REV 03   710-009184                      SIB-M-S
SIB 2            REV 03   710-009184                      SIB-M-S
SIB 3            REV 03   710-009184                      SIB-M-S
Fan Tray 0                                                FFANTRAY-M320-S
Fan Tray 1                                                FFANTRAY-M320-S
Fan Tray 2                                                RFANTRAY-M320-S
```

show chassis hardware (MX5 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                E1368             MX5-T
Midplane         REV 01   711-038215   YF5288            MX5-T
PEM 0            Rev 04   740-028288   VA01215           AC Power Entry Module
PEM 1            Rev 04   740-028288   VA01218           AC Power Entry Module
Routing Engine            BUILTIN      BUILTIN           Routing Engine
TFEB 0                    BUILTIN      BUILTIN           Forwarding Engine
Processor
  QXM 0          REV 05   711-028408   ZA9136            MPC QXM
FPC 0                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN           4x 10GE XFP
    PIC 0                 BUILTIN      BUILTIN           4x 10GE XFP
FPC 1                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0          REV 24   750-028392   YX9820            3D 20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   AM1045SUAQ3       SFP-SX
      Xcvr 1     REV 01   740-031851   AM1045SUAPA       SFP-SX
      Xcvr 2     REV 01   740-031851   AM1045SUAN7       SFP-SX
      Xcvr 3     REV 01   740-031851   AM1045SU91Q       SFP-SX
      Xcvr 4     REV 01   740-031851   AM1045SUDDR       SFP-SX
      Xcvr 9     REV 01   740-011613   AM0848SB6A1       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   AM1045SUAN0       SFP-SX
      Xcvr 1     REV 01   740-011613   AS0812S0719       SFP-SX
      Xcvr 2     REV 01   740-011613   AM0821SA121       SFP-SX
      Xcvr 3     REV 01   740-011613   PF21K21           SFP-SX
      Xcvr 4     REV 01   740-011613   AM0848SB69Z       SFP-SX
      Xcvr 5     REV 01   740-011782   P9P0XV3           SFP-SX
      Xcvr 6     REV 01   740-011613   AM0812S8WJN       SFP-SX
      Xcvr 7     REV 01   740-011613   PAM3G9Q           SFP-SX
      Xcvr 8     REV 01   740-011613   AM0848SB4A6       SFP-SX
      Xcvr 9     REV 01   740-011782   P9M0U37           SFP-SX
  MIC 1          REV 20   750-028380   ZG2657            3D 2x 10GE XFP
    PIC 2                 BUILTIN      BUILTIN           1x 10GE XFP
    PIC 3                 BUILTIN      BUILTIN           1x 10GE XFP
Fan Tray                                                 Fan Tray
```


show chassis hardware (MX10 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                E1372             MX10-T
Midplane         REV 01   711-038211   YF5285            MX10-T
PEM 0            Rev 04   740-028288   VB01678           AC Power Entry Module
Routing Engine            BUILTIN      BUILTIN           Routing Engine
TFEB 0                    BUILTIN      BUILTIN           Forwarding Engine
Processor
  QXM 0          REV 05   711-028408   ZA9053            MPC QXM
FPC 0                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN           4x 10GE XFP
    PIC 0                 BUILTIN      BUILTIN           4x 10GE XFP
FPC 1                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0          REV 24   750-028392   YX9436            3D 20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   AM1107SUFQW       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
Fan Tray                                                 Fan Tray
```

show chassis hardware (MX40 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                E1367             MX40-T
Midplane         REV 01   711-038211   YF5284            MX40-T
PEM 0            Rev 04   740-028288   VB01680           AC Power Entry Module
PEM 1            Rev 04   740-028288   VB01700           AC Power Entry Module
Routing Engine            BUILTIN      BUILTIN           Routing Engine
TFEB 0                    BUILTIN      BUILTIN           Forwarding Engine
Processor
  QXM 0          REV 05   711-028408   ZA9048            MPC QXM
FPC 0                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN           4x 10GE XFP
    PIC 0                 BUILTIN      BUILTIN           4x 10GE XFP
      Xcvr 0     REV 01   740-014279   M7067UPP          XFP-10G-LR
      Xcvr 1              NON-JNPR     K9J02UN           XFP-10G-LR
FPC 1                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0          REV 24   750-028392   YX3504            3D 20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-011613   AM0812S8WTE       SFP-SX
      Xcvr 1     REV 01   740-011613   PFA6KV2           SFP-SX
      Xcvr 2     REV 01   740-031851   AM1045SUDDM       SFP-SX
      Xcvr 3     REV 01   740-011613   PD63C7M           SFP-SX
      Xcvr 4     REV 01   740-011613   PD63DJY           SFP-SX
      Xcvr 5     REV 02   740-011613   AA0950STLL9       SFP-SX
      Xcvr 6     REV 01   740-011782   PAR1YHC           SFP-SX
      Xcvr 7     REV 01   740-011782   P9P0XXL           SFP-SX
      Xcvr 8     REV 01   740-011613   PD63D95           SFP-SX
      Xcvr 9     REV 01   740-031851   AM1045SU9B8       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-011613   PF21L3Z           SFP-SX
      Xcvr 1     REV 01   740-031851   AM1045SU7M9       SFP-SX
      Xcvr 2     REV 01   740-031851   AM1045SUAPT       SFP-SX
      Xcvr 3     REV 01   740-011613   PFF2BZH           SFP-SX
      Xcvr 4     REV 01   740-031851   AM1045SUDDN       SFP-SX
      Xcvr 5     REV 01   740-031851   AM1039S00ZR       SFP-SX
```



```
      Xcvr 6      REV 01   740-031851   AM1045SUD6Y      SFP-SX
      Xcvr 8      REV 01   740-011613   PFM1QBS          SFP-SX
      Xcvr 9      REV 01   740-011613   PFF2E25          SFP-SX
  MIC 1           REV 01   750-021130   KG4391           3D 2x 10GE XFP
    PIC 2                  BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0      REV 01   740-011571   C645XJ04G        XFP-10G-SR
    PIC 3                  BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0               NON-JNPR     CA49BK0AE        XFP-10G-SR
Fan Tray                                                 Fan Tray
```

show chassis hardware (Fixed MX80 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                                  MX80-48T
Midplane          REV 01   711-031603   KF9250           MX80-48T
Routing Engine             BUILTIN      BUILTIN          Routing Engine
FEB 0                      BUILTIN      BUILTIN          Forwarding Engine Board
FPC 0                      BUILTIN      BUILTIN          MPC BUILTIN
  MIC 0                    BUILTIN      BUILTIN          4x 10GE XFP
    PIC 0                  BUILTIN      BUILTIN          4x 10GE XFP
      Xcvr 0               NON-JNPR     M6439D41         XFP-10G-LR
      Xcvr 1      REV 01   740-014279   6XE931N00202     XFP-10G-LR
      Xcvr 2      REV 01   740-014289   C715XU05F        XFP-10G-SR
      Xcvr 3      REV 01   740-014289   C650XU0EP        XFP-10G-SR
FPC 1                      BUILTIN      BUILTIN          MPC BUILTIN
  MIC 0           REV 01   711-029399   JR6981           12x 1GE(LAN) RJ45
    PIC 0                  BUILTIN      BUILTIN          12x 1GE(LAN) RJ45
    PIC 1                  BUILTIN      BUILTIN          12x 1GE(LAN) RJ45
  MIC 1           REV 01   BUILTIN      BUILTIN          12x 1GE(LAN) RJ45
    PIC 2                  BUILTIN      BUILTIN          12x 1GE(LAN) RJ45
    PIC 3                  BUILTIN      BUILTIN          12x 1GE(LAN) RJ45
Fan Tray                                                 Fan Tray
```

show chassis hardware (Modular MX80 Router)

```
user@host> show chassis hardware

Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                                  MX80
Midplane          REV 02   711-031594   JR7084           MX80
PEM 0             Rev 01   740-028288   000018           AC Power Entry Module
Routing Engine             BUILTIN      BUILTIN          Routing Engine
FEB 0                      BUILTIN      BUILTIN          Forwarding Engine Board

  QXM 0           REV 05   711-028408   JR7041           MPC QXM
FPC 0                      BUILTIN      BUILTIN          MPC BUILTIN
  MIC 0                    BUILTIN      BUILTIN          4x 10GE XFP
    PIC 0                  BUILTIN      BUILTIN          4x 10GE XFP
FPC 1                      BUILTIN      BUILTIN          MPC BUILTIN
  MIC 0           REV 02   750-028380   JR6598           3D 2x 10GE XFP
    PIC 0                  BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0      REV 01   740-014289   T07M86365        XFP-10G-SR
    PIC 1                  BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0      REV 01   740-014289   T07M71094        XFP-10G-SR
  MIC 1           REV 02   750-028380   JG8548           3D 2x 10GE XFP
    PIC 2                  BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0      REV 02   740-014289   T08L86302        XFP-10G-SR
```

```
    PIC 3                BUILTIN      BUILTIN           1x 10GE XFP
      Xcvr 0     REV 02   740-014289   C810XU0BA         XFP-10G-SR
Fan Tray                                                 Fan Tray
```

show chassis hardware (MX150)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DD2316AF0078      MX150
Midplane         REV 04   650-066113   DD2316AF0078      MX150
Power Supply 0
Routing Engine 0                                         RE-VMX
CB 0                                                     VMX SCB
CB 1                                                     VMX SCB
FPC 0                                                    Virtual FPC
  CPU            Rev. 1.0 RIOT         BUILTIN
  MIC 0                                                  Virtual
    PIC 0                BUILTIN      BUILTIN           Virtual
      Xcvr 10    REV 02   740-013111   A331846           SFP-T
      Xcvr 11    REV 02   740-013111   C248517           SFP-T
Fan Tray 0                                              fan-ctrl-0 0, Front to
Back Airflow - AFO
Fan Tray 1                                              fan-ctrl-0 1, Front to
Back Airflow - AFO
```

show chassis hardware models (MX150)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 04   650-066113   DD2316AF0163      MX150
Fan Tray 0                                               Assy,Sub,Fan
Tray,AFO,Opus-AFO
Fan Tray 1                                               Assy,Sub,Fan
Tray,AFO,Opus-AFO
```

show chassis hardware (MX104 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                G3503             MX104
Midplane         REV 28   750-044219   CAAX5741          MX104
PEM 0            REV 03   740-045933   1H072500016       AC Power Entry Module
PEM 1            REV 03   740-045932   1H073050017       DC Power Entry Module
Routing Engine 0 REV 20   750-044228   CAAY7935          RE-MX-104
Routing Engine 1 REV 13   750-044228   CAAM6380          RE-MX-104
AFEB 0                    BUILTIN      BUILTIN           Forwarding Engine
Processor
FPC 0                     BUILTIN      BUILTIN           MPC BUILTIN
FPC 1                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0          REV 15   750-036132   CAAF7948          2xOC12/8xOC3 CC-CE
    PIC 0                BUILTIN      BUILTIN           2xOC12/8xOC3 CC-CE
      Xcvr 0     REV 01   740-011615   PCQ0U2J           SFP-IR
      Xcvr 1     REV 01   740-016068   PJL7A6G           SFP-SR
      Xcvr 2     REV 01   740-016068   PJL7A5J           SFP-SR
      Xcvr 3     REV 01   740-016065   PJN5HPZ           SFP-SR
      Xcvr 4     REV 01   740-029122   PKB38TL           SFP-LR
      Xcvr 5     REV 01   740-011787   P6A107G           SFP-LR
```

```
        Xcvr 6      REV 01    740-029122   PKB38TR           SFP-LR
        Xcvr 7      REV 01    740-011787   PBK0NK3           SFP-LR
  MIC 1
FPC 2                         BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                       BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    PIC 0                     BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01    740-031980   B10F00465         SFP+-10G-SR
      Xcvr 1        REV 01    740-031980   B10F00461         SFP+-10G-SR
      Xcvr 2        REV 01    740-031980   B10G01545         SFP+-10G-SR
      Xcvr 3        REV 01    740-031980   B10G01385         SFP+-10G-SR
Fan Tray 0          REV 02    711-049570   CAAX6538          Fan Tray
```

show chassis hardware detail (MX104 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item                Version   Part number  Serial number     Description
Chassis                                    G3503             MX104
Midplane            REV 28    750-044219   CAAX5741          MX104
PEM 0               REV 03    740-045933   1H072500016       AC Power Entry Module
PEM 1               REV 03    740-045932   1H073050017       DC Power Entry Module
Routing Engine 0 REV 20       750-044228   CAAY7935          RE-MX-104
  da0     7836 MB  ATP IG eUSB SSD                           Nand Flash 0
  usb0 (addr 1)  EHCI root hub 0           Freescale         uhub0
  usb0 (addr 2)  USB2513Bi  9491           SMSC              uhub1
  usb0 (addr 3)  ATP IG eUSB SSD 44801 ATP Electronics       umass0
Routing Engine 1 REV 13       750-044228   CAAM6380          RE-MX-104
  da0     7836 MB  ATP IG eUSB SSD                           Nand Flash 0
AFEB 0                        BUILTIN      BUILTIN           Forwarding Engine
Processor
FPC 0                         BUILTIN      BUILTIN           MPC BUILTIN
FPC 1                         BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0             REV 15    750-036132   CAAF7948          2xOC12/8xOC3 CC-CE
    PIC 0                     BUILTIN      BUILTIN           2xOC12/8xOC3 CC-CE
      Xcvr 0        REV 01    740-011615   PCQ0U2J           SFP-IR
      Xcvr 1        REV 01    740-016068   PJL7A6G           SFP-SR
      Xcvr 2        REV 01    740-016068   PJL7A5J           SFP-SR
      Xcvr 3        REV 01    740-016065   PJN5HPZ           SFP-SR
      Xcvr 4        REV 01    740-029122   PKB38TL           SFP-LR
      Xcvr 5        REV 01    740-011787   P6A107G           SFP-LR
      Xcvr 6        REV 01    740-029122   PKB38TR           SFP-LR
      Xcvr 7        REV 01    740-011787   PBK0NK3           SFP-LR
  MIC 1
FPC 2                         BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                       BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    PIC 0                     BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01    740-031980   B10F00465         SFP+-10G-SR
      Xcvr 1        REV 01    740-031980   B10F00461         SFP+-10G-SR
      Xcvr 2        REV 01    740-031980   B10G01545         SFP+-10G-SR
      Xcvr 3        REV 01    740-031980   B10G01385         SFP+-10G-SR
Fan Tray 0          REV 02    711-049570   CAAX6538          Fan Tray
```

show chassis hardware detail (MX480 Packet Transport Router with details of virtual disk size)

```
user@host> show chassis hardware detail
Hardware inventory:
Item                Version   Part number  Serial number     Description
Chassis                                    JN122FFD9AFB      MX480
Midplane            REV 05    710-017414   ACRB8882          MX480 Midplane
FPM Board           REV 02    710-017254   CADF7623          Front Panel Display
```



```
PEM 0            Rev 07   740-017343   QCS1128A0TY       DC Power Entry Module
PEM 1            Rev 07   740-017343   QCS1128A0JM       DC Power Entry Module
Routing Engine 0 REV 07   750-054758   CADG2028          RE-S-2X00x6
  vtbd0 15361 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    511 MB  QEMU HARDDISK        QM00002           Emulated IDE Disk
  usb0 (addr 1)  UHCI root HUB 0       Intel             uhub0
Routing Engine 1 REV 00   750-054758                     RE-S-2X00x6
  vtbd0 15361 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    511 MB  QEMU HARDDISK        QM00002           Emulated IDE Disk
  usb0 (addr 1)  UHCI root HUB 0       Intel             uhub0
CB 0             REV 01   750-055976   CACS1837          Enhanced MX SCB 2
CB 1             REV 01   750-055976   CADD9894          Enhanced MX SCB 2
  Xcvr 1         REV 01   740-031980   AP41KCL           SFP+-10G-SR
FPC 0            REV 09   750-049040   CACX1759          LOAD MPC Type 2
  CPU            REV 10   711-035209   CACP9324          HMPC PMB 2G
FPC 4            REV 28   750-037355   CACY8384          MPC4E 3D 2CGE+8XGE
  CPU            REV 10   711-035209   CACX0428          HMPC PMB 2G
Fan Tray                                                 Enhanced Left Fan Tray
```

show chassis hardware extensive (MX104 Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                G3503             MX104
 Jedec Code:    0x7fb0           EEPROM Version:    0x02
                                 S/N:               G3503
 Assembly ID:  0x0560            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: MX104
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 60 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 47 33 35 30 33 00 00 00 00 00 00 00 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 28   750-044219   CAAX5741          MX104
 Jedec Code:    0x7fb0           EEPROM Version:    0x02
 P/N:          750-044219        S/N:               CAAX5741
 Assembly ID:  0x0560            Assembly Version:  01.28
 Date:         03-27-2013        Assembly Flags:    0x00
 Version:      REV 28            CLEI Code:         PROTOXCLEI
 ID: MX104                       FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ad 01 08 00 b0 a8 6e a7 f8 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 60 01 1c 52 45 56 20 32 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 31 39 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 58 35 37 34 31 00 1b 03 07
  Address 0x30: dd ff ff ff ad 01 08 00 b0 a8 6e a7 f8 00 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 47 33 35 30 33 00 00 00 00 00 00 00
```

```
PEM 0           REV 03   740-045933   1H072500016       AC Power Entry Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-045933       S/N:               1H072500016
 Assembly ID:  0x0475           Assembly Version:  00.03
 Date:         12-14-2012       Assembly Flags:    0x00
 Version:      REV 03           CLEI Code:         IPUPAJ9KAA
 ID: AC Power Entry Module      FRU Model Number:  PWR-AMX1100-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 02 02 00 ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 75 00 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 39 33 33 00 00
  Address 0x20: 31 48 30 37 32 35 30 30 30 31 36 00 00 0e 0c 07
  Address 0x30: dc 30 43 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 02 02 00 ff 01 49 50 55 50 41 4a 39 4b 41 41 50
  Address 0x50: 57 52 2d 41 4d 58 31 31 30 30 2d 41 43 2d 53 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 70 ff ff ff ff ff ff ff ff ff ff ff ff
PEM 1           REV 03   740-045932   1H073050017       DC Power Entry Module
Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-045932       S/N:               1H073050017
 Assembly ID:  0x0476           Assembly Version:  00.03
 Date:         01-30-2013       Assembly Flags:    0x00
 Version:      REV 03           CLEI Code:         IPUPAJ8KAA
 ID: DC Power Entry Module      FRU Model Number:  PWR-AMX1100-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 02 02 00 ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 76 00 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 39 33 32 00 00
  Address 0x20: 31 48 30 37 33 30 35 30 30 31 37 00 00 1e 01 07
  Address 0x30: dd 30 44 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 02 02 00 ff 01 49 50 55 50 41 4a 38 4b 41 41 50
  Address 0x50: 57 52 2d 41 4d 58 31 31 30 30 2d 44 43 2d 53 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 72 ff ff ff ff ff ff ff ff ff ff ff ff
Routing Engine 0 REV 20   750-044228   CAAY7935          RE-MX-104
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-044228       S/N:               CAAY7935
 Assembly ID:  0x0b81           Assembly Version:  01.20
 Date:         03-18-2013       Assembly Flags:    0x00
 Version:      REV 20           CLEI Code:         PROTOXCLEI
 ID: RE-MX-104                  FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ad 01 00 08 b0 a8 6e a6 fc 10 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 81 01 14 52 45 56 20 32 30 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 32 38 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 59 37 39 33 35 00 12 03 07
  Address 0x30: dd ff ff ff ad 01 00 08 b0 a8 6e a6 fc 10 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
  da0    7836 MB  ATP IG eUSB SSD                        Nand Flash 0
  usb0 (addr 1)  EHCI root hub 0       Freescale         uhub0
  usb0 (addr 2)  USB2513Bi  9491       SMSC              uhub1
  usb0 (addr 3)  ATP IG eUSB SSD 44801 ATP Electronics   umass0
Routing Engine 1 REV 13   750-044228   CAAM6380          RE-MX-104
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-044228       S/N:               CAAM6380
```



```
 Assembly ID:  0x0b81           Assembly Version:  01.13
 Date:         09-17-2012       Assembly Flags:    0x00
 Version:      REV 13           CLEI Code:         PROTOXCLEI
 ID: RE-MX-104                  FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ad 01 00 08 64 87 88 27 08 18 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 81 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 32 38 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 4d 36 33 38 30 00 11 09 07
  Address 0x30: dc ff ff ff ad 01 00 08 64 87 88 27 08 18 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
  da0    7836 MB  ATP IG eUSB SSD                        Nand Flash 0
AFEB 0                   BUILTIN      BUILTIN           Forwarding Engine
Processor
FPC 0                    BUILTIN      BUILTIN           MPC BUILTIN
FPC 1                    BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0          REV 15   750-036132   CAAF7948         2xOC12/8xOC3 CC-CE
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-036132       S/N:               CAAF7948
 Assembly ID:  0x0a1a           Assembly Version:  01.15
 Date:         07-03-2012       Assembly Flags:    0x00
 Version:      REV 15           CLEI Code:         IP9IAM2DAA
 ID: 2xOC12/8xOC3 CC-CE         FRU Model Number:  MIC-3D-8OC3-2OC12-ATM
 Board Information Record:
  Address 0x00: 12 01 05 03 05 ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0a 1a 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 36 31 33 32 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 46 37 39 34 38 00 03 07 07
  Address 0x30: dc ff ff ff 12 01 05 03 05 ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 39 49 41 4d 32 44 41 41 4d
  Address 0x50: 49 43 2d 33 44 2d 38 4f 43 33 2d 32 4f 43 31 32
  Address 0x60: 2d 41 54 4d 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff e3 c0 02 a3 9c 00 00 00 00 0a 60 00 00
    PIC 0                BUILTIN      BUILTIN           2xOC12/8xOC3 CC-CE
      Xcvr 0     REV 01   740-011615   PCQ0U2J          SFP-IR
      Xcvr 1     REV 01   740-016068   PJL7A6G          SFP-SR
      Xcvr 2     REV 01   740-016068   PJL7A5J          SFP-SR
      Xcvr 3     REV 01   740-016065   PJN5HPZ          SFP-SR
      Xcvr 4     REV 01   740-029122   PKB38TL          SFP-LR
      Xcvr 5     REV 01   740-011787   P6A107G          SFP-LR
      Xcvr 6     REV 01   740-029122   PKB38TR          SFP-LR
      Xcvr 7     REV 01   740-011787   PBK0NK3          SFP-LR
  MIC 1
FPC 2                    BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
 Jedec Code:   0x0000           EEPROM Version:    0x00
 P/N:          BUILTIN          S/N:               BUILTIN
 Assembly ID:  0x0a60           Assembly Version:  00.00
 Date:         00-00-0000       Assembly Flags:    0x00
 ID: 4x 10GE(LAN) SFP+
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 60 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
  Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
```

```
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 c0 02 a5 04 7f b0 02 ff 0a 1a 01 0f
    PIC 0                 BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   B10F00465         SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   B10F00461         SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   B10G01545         SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   B10G01385         SFP+-10G-SR
Fan Tray 0        REV 02   711-049570   CAAX6538          Fan Tray
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           711-049570        S/N:               CAAX6538
 Assembly ID:   0x0b82            Assembly Version:  01.02
 Date:          03-01-2013        Assembly Flags:    0x00
 Version:       REV 02            CLEI Code:         PROTOXCLEI
 ID: Fan Tray                     FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 82 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 39 35 37 30 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 58 36 35 33 38 00 01 03 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware extensive (PTX10008 Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item              Version  Part number  Serial number     Description
Chassis                                 DE487             JNP10008 [PTX10008 -
PILOT BUILD V1.1]
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
                                  S/N:               DE487
 Assembly ID:   0x0566            Assembly Version:  01.27
 Date:          08-08-2016        Assembly Flags:    0x00
                                  CLEI Code:         CMMUM00ARA
 ID: JNP10008                     FRU Model Number:  QFX10008-CHAS
 Board Information Record:
  Address 0x00: ad 01 08 00 30 b6 4f e9 74 c4 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 66 01 1b 00 45 56 20 32 37 00 00
  Address 0x10: 00 00 00 00 00 35 30 2d 30 35 34 30 39 37 00 00
  Address 0x20: 44 45 34 38 37 00 00 00 00 00 00 00 00 08 08 07
  Address 0x30: e0 ff ff ff ad 01 08 00 30 b6 4f e9 74 c4 ff ff
  Address 0x40: ff ff ff ff 01 43 4d 4d 55 4d 30 30 41 52 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 43 48 41 53 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 44 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 63 44 45 34 38 37 00 00 00 00 00 00 00
Midplane          REV 27   750-054097   ACPD4307          Midplane 8
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           750-054097        S/N:               ACPD4307
 Assembly ID:   0x0be3            Assembly Version:  01.27
 Date:          08-08-2016        Assembly Flags:    0x00
 Version:       REV 27            CLEI Code:         CMMUM00ARA
 ID: QFX10008 Midplane            FRU Model Number:  QFX10008-CHAS
 Board Information Record:
```



```
  Address 0x00: ad 01 08 00 30 b6 4f e9 74 c4 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e3 01 1b 52 45 56 20 32 37 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 34 30 39 37 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 44 34 33 30 37 00 08 08 07
  Address 0x30: e0 ff ff ff ad 01 08 00 30 b6 4f e9 74 c4 ff ff
  Address 0x40: ff ff ff ff 01 43 4d 4d 55 4d 30 30 41 52 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 43 48 41 53 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 44 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 63 44 45 34 38 37 00 00 00 00 00 00 00
Routing Engine 0          BUILTIN      BUILTIN           RE-PTX-2X00x4
  vtbd0 15360 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    128 MB  QEMU                 QM00002           Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0      Intel             uhub0
  usb1 (addr 0.2) product 0x0020 32    vendor 0x8087     uhub1
Routing Engine 1          BUILTIN      BUILTIN           RE-PTX-2X00x4
  vtbd0 15360 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    128 MB  QEMU                 QM00002           Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0      Intel             uhub0
  usb1 (addr 0.2) product 0x0020 32    vendor 0x8087     uhub1
CB 0             REV 02   750-068820   ACNZ4440          Control Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-068820        S/N:               ACNZ4440
 Assembly ID:  0x0b9d            Assembly Version:  01.02
 Date:         06-13-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUCAH3CTB
 ID: Control Board               FRU Model Number:  QFX10000-RE
 Board Information Record:
  Address 0x00: ad 01 00 10 84 c1 c1 54 10 be ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 9d 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 5a 34 34 34 30 00 0d 06 07
  Address 0x30: e0 ff ff ff ad 01 00 10 84 c1 c1 54 10 be ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 33 43 54 42 51
  Address 0x50: 46 58 31 30 30 30 30 2d 52 45 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff db ff ff ff ff ff ff ff ff ff ff ff ff
CB 1             REV 02   750-068820   ACNZ8284          Control Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-068820        S/N:               ACNZ8284
 Assembly ID:  0x0b9d            Assembly Version:  01.02
 Date:         06-27-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUCAH3CTB
 ID: Control Board               FRU Model Number:  QFX10000-RE
 Board Information Record:
  Address 0x00: ad 01 00 10 84 c1 c1 e5 b1 46 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 9d 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 5a 38 32 38 34 00 1b 06 07
  Address 0x30: e0 ff ff ff ad 01 00 10 84 c1 c1 e5 b1 46 ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 33 43 54 42 51
  Address 0x50: 46 58 31 30 30 30 30 2d 52 45 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff db ff ff ff ff ff ff ff ff ff ff ff ff
FPC 0            REV 36   750-051354   ACNP4679          LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
```



```
P/N:          750-051354        S/N:              ACNP4679
Assembly ID:  0x0be7            Assembly Version:  01.36
Date:         11-11-2016        Assembly Flags:    0x00
Version:      REV 36            CLEI Code:         CMUIAM9BAA
ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b e7 01 24 52 45 56 20 33 36 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 35 31 33 35 34 00 00
 Address 0x20: 53 2f 4e 20 41 43 4e 50 34 36 37 39 00 0b 0b 07
 Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 41 51
 Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 42 45 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff fe ff ff ff ff ff ff ff ff ff ff ff ff
 CPU                      BUILTIN      BUILTIN           FPC CPU
Jedec Code:   0x7fb0            EEPROM Version:    0x02
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0xf020            Assembly Version:  02.17
Date:         04-19-2012        Assembly Flags:    0x00
Board Information Record:
 Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff f0 20 02 11 00 e0 3c fa 09 00 70 87
 Address 0x10: 09 38 bb ff 42 55 49 4c 54 49 4e 00 00 e0 3c fa
 Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
 Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
 Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
 PIC 0                  BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

Jedec Code:   0x7fb0            EEPROM Version:    0x02
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0xf050            Assembly Version:  02.17
Date:         04-19-2012        Assembly Flags:    0x00
Board Information Record:
 Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
 Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
 Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
 Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
   Xcvr 1       REV 01   740-058734   1ECQ113834D       QSFP-100GBASE-SR4
   Xcvr 5       REV 01   740-058734   1ECQ1137067       QSFP-100GBASE-SR4
   Xcvr 6       REV 01   740-054053   QF3205SD          QSFP+-4X10G-SR
   Xcvr 7       REV 01   740-058734   1ECQ11381MP       QSFP-100GBASE-SR4
   Xcvr 11      REV 01   740-061405   1ACQ110507K       QSFP-100GBASE-SR4
   Xcvr 13      REV 01   740-058734   1ECQ11390ZB       QSFP-100GBASE-SR4
   Xcvr 17      REV 01   740-058734   1ECQ11381M1       QSFP-100GBASE-SR4
   Xcvr 19      REV 01   740-058734   1ECQ11381JS       QSFP-100GBASE-SR4
   Xcvr 23      REV 01   740-058734   1ACQ112000E       QSFP-100GBASE-SR4
   Xcvr 25      REV 01   740-058734   1ECQ11381NT       QSFP-100GBASE-SR4
   Xcvr 28      REV 01   740-054053   QG1502WV          QSFP+-4X10G-SR
   Xcvr 29      REV 01   740-058734   1ACQ112000D       QSFP-100GBASE-SR4
```



```
    Xcvr 33      REV 01   740-058734   1ACQ1134065       QSFP-100GBASE-SR4
    Xcvr 34      REV 01   740-067442   XV20L4L           QSFP+-40G-SR4
FPC 1            REV 33   750-051354   ACNX8831          LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-051354        S/N:               ACNX8831
 Assembly ID:  0x0be7            Assembly Version:  01.33
 Date:         06-03-2016        Assembly Flags:    0x00
 Version:      REV 33            CLEI Code:         CMUIAM9BAA
 ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e7 01 21 52 45 56 20 33 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 31 33 35 34 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 58 38 38 33 31 00 03 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff fb ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                        BUILTIN      BUILTIN           FPC CPU
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 20 3e fa 09 00 10 8a
  Address 0x10: 09 38 bb ff 42 55 49 4c 54 49 4e 00 00 20 3e fa
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                    BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf050            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
    Xcvr 5                   NON-JNPR     37700171YY0084    QSFP-100GBASE-LR4
    Xcvr 25                  NON-JNPR     GDA2017459        QSFP-100GBASE-LR4
    Xcvr 29                  NON-JNPR     GDF2008750        QSFP-100GBASE-LR4
FPC 2           REV 32   750-051357   ACPB0341         LC1101 - 30C / 30Q / 96X

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-051357        S/N:               ACPB0341
```

```
 Assembly ID:  0x0be8           Assembly Version:  01.32
 Date:         06-04-2016       Assembly Flags:    0x00
 Version:      REV 32           CLEI Code:         CMUIANABAA
 ID: ULC-30Q28                  FRU Model Number:  QFX10000-30C
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e8 01 20 52 45 56 20 33 32 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 31 33 35 37 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 42 30 33 34 31 00 04 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4e 41 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 30 43 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ef ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                     BUILTIN      BUILTIN           FPC CPU
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          BUILTIN          S/N:               BUILTIN
 Assembly ID:  0xf020           Assembly Version:  02.17
 Date:         04-19-2012       Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 00 67 00 0a 00 b0 8c
  Address 0x10: 09 38 bb ff 42 55 49 4c 54 49 4e 00 00 00 67 00
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                  BUILTIN      BUILTIN           30x100GE/30x40GE/96x10GE

 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          BUILTIN          S/N:               BUILTIN
 Assembly ID:  0xf050           Assembly Version:  02.17
 Date:         04-19-2012       Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 de ad be ef de ad be ef de ad be ef
    Xcvr 0                NON-JNPR     37700170YZC305    QSFP-100GBASE-LR4
    Xcvr 4                NON-JNPR     37700170YZC306    QSFP-100GBASE-LR4
    Xcvr 9       REV 01   740-054053   QF36013S          QSFP+-4X10G-SR
    Xcvr 12      REV 01   740-067442   XV301AU           QSFP+-40G-SR4
    Xcvr 14      REV 01   740-043308   UWE2CG9           QSFP+-40G-LR4
    Xcvr 16      REV 01   740-043308   UWH141S           QSFP+-40G-LR4
    Xcvr 17      REV 01   740-058734   1ECQ11180VH       QSFP-100GBASE-SR4
    Xcvr 18      REV 01   740-054050   INFAJ0492237      QSFP+-4X10G-LR
    Xcvr 26      REV 01   740-058734   1ACQ111803N       QSFP-100GBASE-SR4
    Xcvr 27      REV 01   740-058734   1ACQ113405S       QSFP-100GBASE-SR4
FPC 3            REV 35   750-051357   ACPD2186         LC1101 - 30C / 30Q / 96X

 Jedec Code:   0x7fb0           EEPROM Version:    0x02
```



```
 P/N:          750-051357        S/N:              ACPD2186
 Assembly ID:  0x0be8            Assembly Version:  01.35
 Date:         09-21-2016        Assembly Flags:    0x00
 Version:      REV 35            CLEI Code:         CMUIANABAA
 ID: ULC-30Q28                   FRU Model Number:  QFX10000-30C
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e8 01 23 52 45 56 20 33 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 31 33 35 37 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 44 32 31 38 36 00 15 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4e 41 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 30 43 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 44 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f1 ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                    BUILTIN      BUILTIN            FPC CPU
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 80 70 fa 09 00 50 8f
  Address 0x10: 09 38 bb ff 42 55 49 4c 54 49 4e 00 00 80 70 fa
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                 BUILTIN      BUILTIN           30x100GE/30x40GE/96x10GE

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf050            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
    Xcvr 0       REV 01   740-061409   1GCQA1470A3       QSFP-100GBASE-LR4
    Xcvr 1       REV 01   740-061409   1GCQA1470XC       QSFP-100GBASE-LR4
    Xcvr 7                NON-JNPR     FG4550500008      QSFP-100G-CWDM4
    Xcvr 24      REV 01   740-058734   1ECQ11381LX       QSFP-100GBASE-SR4
    Xcvr 29      REV 01   740-043308   UWE0UYS           QSFP+-40G-LR4
FPC 5            REV 08   750-068822   ACPF0057          LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-068822        S/N:               ACPF0057
 Assembly ID:  0x0be7            Assembly Version:  01.08
 Date:         09-01-2016        Assembly Flags:    0x00
 Version:      REV 08            CLEI Code:         CMUIAM9BAB
```

```
  ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b e7 01 08 52 45 56 20 30 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 32 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 46 30 30 35 37 00 01 09 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 42 51
   Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 42 45 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   CPU                       BUILTIN      BUILTIN           FPC CPU
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          BUILTIN           S/N:               BUILTIN
  Assembly ID:  0xf020            Assembly Version:  02.17
  Date:         04-19-2012        Assembly Flags:    0x00
  Board Information Record:
   Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff f0 20 02 11 00 00 3d fa 09 00 90 94
   Address 0x10: 09 38 bb ff 42 55 49 4c 54 49 4e 00 00 00 3d fa
   Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
   Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
   Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
   PIC 0                    BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          BUILTIN           S/N:               BUILTIN
  Assembly ID:  0xf050            Assembly Version:  02.17
  Date:         04-19-2012        Assembly Flags:    0x00
  Board Information Record:
   Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
   Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
   Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
   Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
   Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
FPC 6            REV 08   750-068822   ACPE9951          LC1102 - 12C / 36Q /
144X
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-068822        S/N:               ACPE9951
  Assembly ID:  0x0be7            Assembly Version:  01.08
  Date:         09-01-2016        Assembly Flags:    0x00
  Version:      REV 08            CLEI Code:         CMUIAM9BAB
  ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b e7 01 08 52 45 56 20 30 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 32 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 45 39 39 35 31 00 01 09 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 42 51
```



```
 Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 42 45 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 CPU                    BUILTIN      BUILTIN           FPC CPU
Jedec Code:   0x7fb0           EEPROM Version:    0x02
P/N:          BUILTIN          S/N:               BUILTIN
Assembly ID:  0xf020           Assembly Version:  02.17
Date:         04-19-2012       Assembly Flags:    0x00
Board Information Record:
 Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff f0 20 02 11 00 c0 3e fa 09 00 30 97
 Address 0x10: 09 38 bb ff 42 55 49 4c 54 49 4e 00 00 c0 3e fa
 Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
 Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
 Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
 PIC 0                 BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

Jedec Code:   0x7fb0           EEPROM Version:    0x02
P/N:          BUILTIN          S/N:               BUILTIN
Assembly ID:  0xf050           Assembly Version:  02.17
Date:         04-19-2012       Assembly Flags:    0x00
Board Information Record:
 Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
 Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
 Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
 Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
   Xcvr 1       REV 01   740-054053   QF3208LG          QSFP+-4X10G-SR
   Xcvr 7       REV 01   740-067442   XV20LGN           QSFP+-40G-SR4
   Xcvr 8       REV 01   740-067442   XV20VMV           QSFP+-40G-SR4
   Xcvr 9       REV 01   740-067442   XV20KCN           QSFP+-40G-SR4
   Xcvr 10      REV 01   740-067442   XU504QD           QSFP+-40G-SR4
   Xcvr 11      REV 01   740-067442   XU504X7           QSFP+-40G-SR4
   Xcvr 12      REV 01   740-067442   XU504W8           QSFP+-40G-SR4
   Xcvr 16      REV 01   740-032986   QF4301JP          QSFP+-40G-SR4
   Xcvr 17      REV 01   740-032986   QF4303AE          QSFP+-40G-SR4
   Xcvr 18      REV 01   740-054050   INFAJ0492400      QSFP+-4X10G-LR
   Xcvr 19      REV 01   740-054050   INFAJ0492142      QSFP+-4X10G-LR
   Xcvr 24      REV 01   740-032986   QF4301KB          QSFP+-40G-SR4
   Xcvr 25      REV 01   740-032986   QF4303YP          QSFP+-40G-SR4
   Xcvr 30      REV 01   740-067442   XV300ZX           QSFP+-40G-SR4
   Xcvr 31      REV 01   740-043308   UWH2KBW           QSFP+-40G-LR4
   Xcvr 34      REV 01   740-054053   QG1501YU          QSFP+-4X10G-SR
FPD Board       REV 07   711-054687   ACPC7142          Front Panel Display
Jedec Code:   0x7fb0           EEPROM Version:    0x01
P/N:          711-054687       S/N:               ACPC7142
Assembly ID:  0x0bf2           Assembly Version:  01.07
Date:         07-22-2016       Assembly Flags:    0x00
Version:      REV 07
ID: QFX10000 FPD
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

```
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b f2 01 07 52 45 56 20 30 37 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 35 34 36 38 37 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 43 37 31 34 32 00 16 07 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 0   REV 02   740-049388   1EDL62102N9       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL62102N9
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         05-25-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 32 31 30 32 4e 39 00 00 19 05 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 1   REV 02   740-049388   1EDL60300KX       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL60300KX
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         01-20-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 30 33 30 30 4b 58 00 00 14 01 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 2   REV 02   740-049388   1EDL60300DL       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL60300DL
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         01-20-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 30 33 30 30 44 4c 00 00 14 01 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
```



```
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 3   REV 02   740-049388   1EDL61701BT       Power Supply AC
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-049388       S/N:               1EDL61701BT
 Assembly ID:  0x0483           Assembly Version:  01.02
 Date:         05-01-2016       Assembly Flags:    0x00
 Version:      REV 02           CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 31 37 30 31 42 54 00 00 01 05 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 4   REV 02   740-049388   1EDL62102P7       Power Supply AC
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-049388       S/N:               1EDL62102P7
 Assembly ID:  0x0483           Assembly Version:  01.02
 Date:         05-25-2016       Assembly Flags:    0x00
 Version:      REV 02           CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 32 31 30 32 50 37 00 00 19 05 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 5   REV 02   740-049388   1EDL62102PP       Power Supply AC
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-049388       S/N:               1EDL62102PP
 Assembly ID:  0x0483           Assembly Version:  01.02
 Date:         05-25-2016       Assembly Flags:    0x00
 Version:      REV 02           CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 32 31 30 32 50 50 00 00 19 05 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
FTC 0            REV 14   750-050108   ACPE4038          Fan Controller 8
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-050108       S/N:               ACPE4038
 Assembly ID:  0x0bee           Assembly Version:  01.14
 Date:         09-27-2016       Assembly Flags:    0x00
```

```
 Version:      REV 14           CLEI Code:         CMUCAHZCAA
 ID: QFX10000 FTC               FRU Model Number:  QFX10008-FAN-CTRL
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ee 01 0e 52 45 56 20 31 34 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 31 30 38 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 45 34 30 33 38 00 1b 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 5a 43 41 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 46 41 4e 2d 43 54 52 4c
  Address 0x60: 00 00 00 00 00 00 41 44 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 98 ff ff ff ff ff ff ff ff ff ff ff ff
FTC 1            REV 14   750-050108   ACPE4032          Fan Controller 8
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-050108       S/N:               ACPE4032
 Assembly ID:  0x0bee           Assembly Version:  01.14
 Date:         09-27-2016       Assembly Flags:    0x00
 Version:      REV 14           CLEI Code:         CMUCAHZCAA
 ID: QFX10000 FTC               FRU Model Number:  QFX10008-FAN-CTRL
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ee 01 0e 52 45 56 20 31 34 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 31 30 38 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 45 34 30 33 32 00 1b 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 5a 43 41 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 46 41 4e 2d 43 54 52 4c
  Address 0x60: 00 00 00 00 00 00 41 44 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 98 ff ff ff ff ff ff ff ff ff ff ff ff
Fan Tray 0       REV 09   760-054372   ACPD6799          Fan Tray 8
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          760-054372       S/N:               ACPD6799
 Assembly ID:  0x0bf0           Assembly Version:  01.09
 Date:         09-28-2016       Assembly Flags:    0x00
 Version:      REV 09           CLEI Code:         CMUCAHYCAA
 ID: QFX10008 FHB               FRU Model Number:  QFX10008-FAN
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b f0 01 09 52 45 56 20 30 39 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 35 34 33 37 32 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 44 36 37 39 39 00 1c 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 59 43 41 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 46 41 4e 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f1 ff ff ff ff ff ff ff ff ff ff ff ff
Fan Tray 1       REV 09   760-054372   ACNZ3584          Fan Tray 8
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          760-054372       S/N:               ACNZ3584
 Assembly ID:  0x0bf0           Assembly Version:  01.09
 Date:         08-30-2016       Assembly Flags:    0x00
 Version:      REV 09           CLEI Code:         CMUCAHYCAA
 ID: QFX10008 FHB               FRU Model Number:  QFX10008-FAN
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b f0 01 09 52 45 56 20 30 39 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 35 34 33 37 32 00 00
```



```
    Address 0x20: 53 2f 4e 20 41 43 4e 5a 33 35 38 34 00 1e 08 07
    Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 59 43 41 41 51
    Address 0x50: 46 58 31 30 30 30 38 2d 46 41 4e 00 00 00 00 00
    Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff f1 ff ff ff ff ff ff ff ff ff ff ff ff
SIB 0            REV 24   750-050058   ACPD4587         Switch Fabric 8
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-050058       S/N:               ACPD4587
  Assembly ID:  0x0bec           Assembly Version:  01.24
  Date:         06-19-2016       Assembly Flags:    0x00
  Version:      REV 24           CLEI Code:         CMUCAH0CAA
  ID: QFX10008 SIB               FRU Model Number:  QFX10008-SF
  Board Information Record:
    Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 0b ec 01 18 52 45 56 20 32 34 00 00
    Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 35 38 00 00
    Address 0x20: 53 2f 4e 20 41 43 50 44 34 35 38 37 00 13 06 07
    Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 30 43 41 41 51
    Address 0x50: 46 58 31 30 30 30 38 2d 53 46 00 00 00 00 00 00
    Address 0x60: 00 00 00 00 00 00 41 45 00 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff d1 00 00 00 00 00 00 00 00 00 00 00 00
SIB 1            REV 24   750-050058   ACNZ0635         Switch Fabric 8
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-050058       S/N:               ACNZ0635
  Assembly ID:  0x0bec           Assembly Version:  01.24
  Date:         06-06-2016       Assembly Flags:    0x00
  Version:      REV 24           CLEI Code:         CMUCAH0CAA
  ID: QFX10008 SIB               FRU Model Number:  QFX10008-SF
  Board Information Record:
    Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 0b ec 01 18 52 45 56 20 32 34 00 00
    Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 35 38 00 00
    Address 0x20: 53 2f 4e 20 41 43 4e 5a 30 36 33 35 00 06 06 07
    Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 30 43 41 41 51
    Address 0x50: 46 58 31 30 30 30 38 2d 53 46 00 00 00 00 00 00
    Address 0x60: 00 00 00 00 00 00 41 45 00 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff d1 00 00 00 00 00 00 00 00 00 00 00 00
SIB 2            REV 24   750-050058   ACPD4908         Switch Fabric 8
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-050058       S/N:               ACPD4908
  Assembly ID:  0x0bec           Assembly Version:  01.24
  Date:         07-12-2016       Assembly Flags:    0x00
  Version:      REV 24           CLEI Code:         CMUCAH0CAA
  ID: QFX10008 SIB               FRU Model Number:  QFX10008-SF
  Board Information Record:
    Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
    Address 0x00: 7f b0 02 ff 0b ec 01 18 52 45 56 20 32 34 00 00
    Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 35 38 00 00
    Address 0x20: 53 2f 4e 20 41 43 50 44 34 39 30 38 00 0c 07 07
    Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
    Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 30 43 41 41 51
    Address 0x50: 46 58 31 30 30 30 38 2d 53 46 00 00 00 00 00 00
    Address 0x60: 00 00 00 00 00 00 41 45 00 ff ff ff ff ff ff ff
    Address 0x70: ff ff ff d1 00 00 00 00 00 00 00 00 00 00 00 00
SIB 3            REV 24   750-050058   ACNZ0617         Switch Fabric 8
```



```
 Jedec Code:   0x7fb0          EEPROM Version:    0x02
 P/N:          750-050058      S/N:               ACNZ0617
 Assembly ID:  0x0bec          Assembly Version:  01.24
 Date:         06-07-2016      Assembly Flags:    0x00
 Version:      REV 24          CLEI Code:         CMUCAH0CAA
 ID: QFX10008 SIB              FRU Model Number:  QFX10008-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ec 01 18 52 45 56 20 32 34 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 35 38 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 5a 30 36 31 37 00 07 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 30 43 41 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 45 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d1 00 00 00 00 00 00 00 00 00 00 00 00
SIB 4            REV 24   750-050058   ACNZ0527         Switch Fabric 8
 Jedec Code:   0x7fb0          EEPROM Version:    0x02
 P/N:          750-050058      S/N:               ACNZ0527
 Assembly ID:  0x0bec          Assembly Version:  01.24
 Date:         06-06-2016      Assembly Flags:    0x00
 Version:      REV 24          CLEI Code:         CMUCAH0CAA
 ID: QFX10008 SIB              FRU Model Number:  QFX10008-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ec 01 18 52 45 56 20 32 34 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 35 38 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 5a 30 35 32 37 00 06 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 30 43 41 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 45 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d1 00 00 00 00 00 00 00 00 00 00 00 00
SIB 5            REV 23   750-050058   ACNX6980         Switch Fabric 8
 Jedec Code:   0x7fb0          EEPROM Version:    0x02
 P/N:          750-050058      S/N:               ACNX6980
 Assembly ID:  0x0bec          Assembly Version:  01.23
 Date:         05-16-2016      Assembly Flags:    0x00
 Version:      REV 23          CLEI Code:         CMUCAH0CAA
 ID: QFX10008 SIB              FRU Model Number:  QFX10008-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ec 01 17 52 45 56 20 32 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 35 38 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 58 36 39 38 30 00 10 05 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 30 43 41 41 51
  Address 0x50: 46 58 31 30 30 30 38 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ce 00 00 00 00 00 00 00 00 00 00 00 00
```

show chassis hardware extensive (PTX10016 Router)

```
user@host> show chassis hardware extensive

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DH995             JNP10016 [PTX10016]
```



```
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
                                 S/N:               DH995
  Assembly ID:  0x0566           Assembly Version:  01.22
  Date:         02-16-2017       Assembly Flags:    0x00
                                 CLEI Code:         CMMUN00ARA
  ID: JNP10016                   FRU Model Number:  QFX10016-CHAS
  Board Information Record:
   Address 0x00: ad 01 10 00 44 aa 50 ab 1b b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 05 66 01 16 00 45 56 20 32 32 00 00
   Address 0x10: 00 00 00 00 00 35 30 2d 30 35 36 35 35 35 00 00
   Address 0x20: 44 48 39 39 35 00 00 00 00 00 00 00 00 10 02 07
   Address 0x30: e1 ff ff ff ad 01 10 00 44 aa 50 ab 1b b6 ff ff
   Address 0x40: ff ff ff ff 01 43 4d 4d 55 4e 30 30 41 52 41 51
   Address 0x50: 46 58 31 30 30 31 36 2d 43 48 41 53 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 32 41 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 51 44 48 39 39 35 00 00 00 00 00 00 00
Midplane         REV 22   750-056555   ACPM7810         Midplane 16
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-056555       S/N:               ACPM7810
  Assembly ID:  0x0be4           Assembly Version:  01.22
  Date:         02-16-2017       Assembly Flags:    0x00
  Version:      REV 22           CLEI Code:         CMMUN00ARA
  ID: QFX10016 Midplane          FRU Model Number:  QFX10016-CHAS
  Board Information Record:
   Address 0x00: ad 01 10 00 44 aa 50 ab 1b b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b e4 01 16 52 45 56 20 32 32 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 35 36 35 35 35 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 4d 37 38 31 30 00 10 02 07
   Address 0x30: e1 ff ff ff ad 01 10 00 44 aa 50 ab 1b b6 ff ff
   Address 0x40: ff ff ff ff 01 43 4d 4d 55 4e 30 30 41 52 41 51
   Address 0x50: 46 58 31 30 30 31 36 2d 43 48 41 53 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 32 41 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 51 44 48 39 39 35 00 00 00 00 00 00 00
Routing Engine 0          BUILTIN      BUILTIN          RE-PTX-2X00x4
  vtbd0 15360 MB                                        Virtio Block Disk
  vtbd1 15360 MB                                        Virtio Block Disk
  ada0    128 MB  QEMU                 QM00002          Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0      Intel            uhub0
  usb1 (addr 0.2) product 0x0020 32    vendor 0x8087    uhub1
Routing Engine 1          BUILTIN      BUILTIN          RE-PTX-2X00x4
  vtbd0 15360 MB                                        Virtio Block Disk
  vtbd1 15360 MB                                        Virtio Block Disk
  ada0    128 MB  QEMU                 QM00002          Virtio Block Disk
  usb0 (addr 0.1) EHCI root HUB 0      Intel            uhub0
  usb1 (addr 0.2) product 0x0020 32    vendor 0x8087    uhub1
CB 0             REV 03   750-068820   ACPL7238         Control Board
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-068820       S/N:               ACPL7238
  Assembly ID:  0x0b9d           Assembly Version:  01.03
  Date:         03-15-2017       Assembly Flags:    0x00
  Version:      REV 03           CLEI Code:         CMUCAH3CTB
  ID: Control Board              FRU Model Number:  QFX10000-RE
  Board Information Record:
   Address 0x00: ad 01 00 10 e8 b6 c2 46 aa 29 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 9d 01 03 52 45 56 20 30 33 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 30 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 4c 37 32 33 38 00 0f 03 07
   Address 0x30: e1 ff ff ff ad 01 00 10 e8 b6 c2 46 aa 29 ff ff
```

```
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 33 43 54 42 51
  Address 0x50: 46 58 31 30 30 30 30 2d 52 45 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff db ff ff ff ff ff ff ff ff ff ff ff ff
CB 1             REV 03   750-068820   ACPL7298         Control Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-068820        S/N:               ACPL7298
 Assembly ID:  0x0b9d            Assembly Version:  01.03
 Date:         03-15-2017        Assembly Flags:    0x00
 Version:      REV 03            CLEI Code:         CMUCAH3CTB
 ID: Control Board               FRU Model Number:  QFX10000-RE
 Board Information Record:
  Address 0x00: ad 01 00 10 e8 b6 c2 46 99 b9 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 9d 01 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4c 37 32 39 38 00 0f 03 07
  Address 0x30: e1 ff ff ff ad 01 00 10 e8 b6 c2 46 99 b9 ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 33 43 54 42 51
  Address 0x50: 46 58 31 30 30 30 30 2d 52 45 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff db ff ff ff ff ff ff ff ff ff ff ff ff
FPC 1            REV 36   750-077140   ACNP4590         LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-077140        S/N:               ACNP4590
 Assembly ID:  0x0be7            Assembly Version:  01.36
 Date:         10-17-2016        Assembly Flags:    0x00
 Version:      REV 36            CLEI Code:         CMUIAM9BAA
 ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e7 01 24 52 45 56 20 33 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 37 37 31 34 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 50 34 35 39 30 00 11 0a 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 45 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff fe ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                     BUILTIN      BUILTIN           FPC CPU
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 40 36 bd 09 40 25 32
  Address 0x10: 09 e8 ba ff 42 55 49 4c 54 49 4e 00 00 40 36 bd
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                 BUILTIN      BUILTIN         12x100GE/36x40GE/144x10GE

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
```



```
  Assembly ID:  0xf050           Assembly Version:  02.17
  Date:         04-19-2012       Assembly Flags:    0x00
  Board Information Record:
   Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
   Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
   Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
   Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
   Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
     Xcvr 0       REV 01   740-054053   QF3600AV         QSFP+-4X10G-SR
     Xcvr 35      REV 01   740-061405   1ACQ110507K      QSFP-100GBASE-SR4
FPC 3            REV 07   750-071975   CAHA2224         LC1102 - 12C / 36Q /
144X
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-071975       S/N:               CAHA2224
  Assembly ID:  0x0be7           Assembly Version:  01.07
  Date:         01-17-2017       Assembly Flags:    0x00
  Version:      REV 07           CLEI Code:         PROTOXCLEI
  ID: ULC-36Q-12Q28              FRU Model Number:  PROTO-ASSEMBLY
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b e7 01 07 52 45 56 20 30 37 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 37 31 39 37 35 00 00
   Address 0x20: 53 2f 4e 20 43 41 48 41 32 32 32 34 00 11 01 07
   Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
   Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
   Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
   CPU                    BUILTIN      BUILTIN           FPC CPU
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          BUILTIN          S/N:               BUILTIN
  Assembly ID:  0xf020           Assembly Version:  02.17
  Date:         04-19-2012       Assembly Flags:    0x00
  Board Information Record:
   Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff f0 20 02 11 00 60 b6 be 09 c0 cf 38
   Address 0x10: 09 e8 ba ff 42 55 49 4c 54 49 4e 00 00 60 b6 be
   Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
   Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
   Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
   PIC 0                 BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          BUILTIN          S/N:               BUILTIN
  Assembly ID:  0xf050           Assembly Version:  02.17
  Date:         04-19-2012       Assembly Flags:    0x00
  Board Information Record:
   Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
   Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
```

```
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
    Xcvr 0       REV 01   740-054053   QG1505YM          QSFP+-4X10G-SR
    Xcvr 11               NON-JNPR     GDA2017459        QSFP-100GBASE-LR4
    Xcvr 35               NON-JNPR     GDF2008750        QSFP-100GBASE-LR4
FPC 5            REV 13   750-068822   ACPD6501          LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-068822        S/N:               ACPD6501
 Assembly ID:  0x0be7            Assembly Version:  01.13
 Date:         06-29-2017        Assembly Flags:    0x00
 Version:      REV 13            CLEI Code:         CMUIAM9BAC
 ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e7 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 38 38 32 32 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 44 36 35 30 31 00 1d 06 07
  Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 43 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 43 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff fd ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                     BUILTIN      BUILTIN           FPC CPU
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 c0 c6 bc 09 c0 ca 40
  Address 0x10: 09 e8 ba ff 42 55 49 4c 54 49 4e 00 00 c0 c6 bc
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                  BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf050            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
```



```
   Xcvr 1       REV 01   740-058734   1ECQ11381LA       QSFP-100GBASE-SR4
   Xcvr 2       REV 01   740-043308   UWH141S           QSFP+-40G-LR4
   Xcvr 3       REV 01   740-043308   UWE2CG9           QSFP+-40G-LR4
FPC 6           REV 37   750-077140   ACNS2793          LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-077140        S/N:               ACNS2793
 Assembly ID:  0x0be7            Assembly Version:  01.37
 Date:         03-25-2017        Assembly Flags:    0x00
 Version:      REV 37            CLEI Code:         CMUIAM9BAA
 ID: ULC-36Q-12Q28               FRU Model Number:  QFX10000-36Q
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e7 01 25 52 45 56 20 33 37 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 37 37 31 34 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 53 32 37 39 33 00 19 03 07
  Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 45 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff fe ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                     BUILTIN      BUILTIN           FPC CPU
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 a0 e6 d4 09 00 bd 43
  Address 0x10: 09 e8 ba ff 42 55 49 4c 54 49 4e 00 00 a0 e6 d4
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                  BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf050            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
   Xcvr 0       REV 01   740-032986   QH0400VH          QSFP+-40G-SR4
   Xcvr 1       REV 01   740-032986   QH0400VM          QSFP+-40G-SR4
   Xcvr 35      REV 01   740-058734   1ECQ11390ZB       QSFP-100GBASE-SR4
FPC 8           REV 36   750-077140   ACNP4625          LC1102 - 12C / 36Q /
144X
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
```

```
 P/N:           750-077140        S/N:               ACNP4625
 Assembly ID:  0x0be7            Assembly Version:  01.36
 Date:          10-17-2016        Assembly Flags:    0x00
 Version:       REV 36            CLEI Code:         CMUIAM9BAA
 ID: ULC-36Q-12Q28                FRU Model Number:  QFX10000-36Q
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e7 01 24 52 45 56 20 33 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 37 37 31 34 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 4e 50 34 36 32 35 00 11 0a 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4d 39 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 33 36 51 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 42 45 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff fe ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                    BUILTIN      BUILTIN           FPC CPU
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:          04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 c0 e6 d4 09 40 59 4a
  Address 0x10: 09 e8 ba ff 42 55 49 4c 54 49 4e 00 00 c0 e6 d4
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                BUILTIN      BUILTIN          12x100GE/36x40GE/144x10GE

 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf050            Assembly Version:  02.17
 Date:          04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
    Xcvr 1       REV 01   740-058732   1AMQA14206D       QSFP-100GBASE-LR4
    Xcvr 10      REV 01   740-032986   QF4301KB          QSFP+-40G-SR4
    Xcvr 24      REV 01   740-054050   INFAJ0492244      QSFP+-4X10G-LR
FPC 9            REV 35   750-071976   ACPD3055          LC1101 - 30C / 30Q / 96X

 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           750-071976        S/N:               ACPD3055
 Assembly ID:  0x0be8            Assembly Version:  01.35
 Date:          05-26-2016        Assembly Flags:    0x00
 Version:       REV 35            CLEI Code:         CMUIANABAA
 ID: ULC-30Q28                    FRU Model Number:  JNP10K-LC1101
 Board Information Record:
```



```
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b e8 01 23 52 45 56 20 33 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 37 31 39 37 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 44 33 30 35 35 00 1a 05 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 49 41 4e 41 42 41 41 4a
  Address 0x50: 4e 50 31 30 4b 2d 4c 43 31 31 30 31 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ef ff ff ff ff ff ff ff ff ff ff ff ff
  CPU                     BUILTIN      BUILTIN           FPC CPU
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf020            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 20 02 11 00 20 e7 d4 09 00 a6 4d
  Address 0x10: 09 e8 ba ff 42 55 49 4c 54 49 4e 00 00 20 e7 d4
  Address 0x20: 42 55 49 4c 54 49 4e 00 42 55 49 4c 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 50 36 36 36 36 00 00 00 00 00 00 00
  PIC 0                  BUILTIN      BUILTIN          30x100GE/30x40GE/96x10GE

 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0xf050            Assembly Version:  02.17
 Date:         04-19-2012        Assembly Flags:    0x00
 Board Information Record:
  Address 0x00: ad 01 01 04 ac 4b c8 1d f7 b6 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff f0 50 02 11 00 00 00 00 07 0a 20 45
  Address 0x10: 6c 61 70 73 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 13 04 07
  Address 0x30: dc ff ff ff ad 01 01 04 ac 4b c8 1d f7 b6 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 45 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f3 55 55 55 55 55 55 55 55 55 55 55 55
    Xcvr 0                NON-JNPR     INGBT7970007      QSFP-100GBASE-LR4
    Xcvr 1                NON-JNPR     UWQ24D9           QSFP-100GBASE-LR4
    Xcvr 2                NON-JNPR     INGBT7970011      QSFP-100GBASE-LR4
    Xcvr 3                NON-JNPR     UX60AF1           QSFP-100G-CWDM4
    Xcvr 4                NON-JNPR     UX408JJ           QSFP-100GBASE-LR4
    Xcvr 11      REV 01   740-058734   1ECQ113835F       QSFP-100GBASE-SR4
    Xcvr 18               NON-JNPR     Q7496             QSFP-100G-CWDM4
    Xcvr 29      REV 01   740-058734   1ECQ11380LZ       QSFP-100GBASE-SR4
Power Supply 0   REV 02   740-049388   1EDL625039E       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL625039E
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         06-19-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
```

```
   Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
   Address 0x20: 31 45 44 4c 36 32 35 30 33 39 45 00 00 13 06 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
   Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
   Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
 Power Supply 1   REV 02   740-049388   1EDL62503AD       Power Supply AC
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          740-049388       S/N:               1EDL62503AD
  Assembly ID:  0x0483           Assembly Version:  01.02
  Date:         06-19-2016       Assembly Flags:    0x00
  Version:      REV 02           CLEI Code:         CMUPADNBAA
  ID: QFX10000 AC                FRU Model Number:  QFX10000-PWR-AC
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
   Address 0x20: 31 45 44 4c 36 32 35 30 33 41 44 00 00 13 06 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
   Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
   Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
 Power Supply 2   REV 02   740-049388   1EDL625039P       Power Supply AC
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          740-049388       S/N:               1EDL625039P
  Assembly ID:  0x0483           Assembly Version:  01.02
  Date:         06-19-2016       Assembly Flags:    0x00
  Version:      REV 02           CLEI Code:         CMUPADNBAA
  ID: QFX10000 AC                FRU Model Number:  QFX10000-PWR-AC
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
   Address 0x20: 31 45 44 4c 36 32 35 30 33 39 50 00 00 13 06 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
   Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
   Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
 Power Supply 3   REV 02   740-049388   1EDL702004E       Power Supply AC
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          740-049388       S/N:               1EDL702004E
  Assembly ID:  0x0483           Assembly Version:  01.02
  Date:         01-18-2017       Assembly Flags:    0x00
  Version:      REV 02           CLEI Code:         CMUPADNBAA
  ID: QFX10000 AC                FRU Model Number:  QFX10000-PWR-AC
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
   Address 0x20: 31 45 44 4c 37 30 32 30 30 34 45 00 00 12 01 07
   Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
   Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
   Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
```

```
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 4   REV 02   740-049388   1EDL625039D       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL625039D
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         06-19-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 32 35 30 33 39 44 00 00 13 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 5   REV 02   740-049388   1EDL63706JD       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL63706JD
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         09-13-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 33 37 30 36 4a 44 00 00 0d 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
Power Supply 6   REV 02   740-049388   1EDL63706JH       Power Supply AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-049388        S/N:               1EDL63706JH
 Assembly ID:  0x0483            Assembly Version:  01.02
 Date:         09-13-2016        Assembly Flags:    0x00
 Version:      REV 02            CLEI Code:         CMUPADNBAA
 ID: QFX10000 AC                 FRU Model Number:  QFX10000-PWR-AC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 83 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 39 33 38 38 00 00
  Address 0x20: 31 45 44 4c 36 33 37 30 36 4a 48 00 00 0d 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 50 41 44 4e 42 41 41 51
  Address 0x50: 46 58 31 30 30 30 30 2d 50 57 52 2d 41 43 00 00
  Address 0x60: 00 00 00 00 00 00 01 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff dc ff ff ff ff ff ff ff ff ff ff ff ff
FTC 0            REV 10   750-050309   ACPM2918          Fan Controller 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-050309        S/N:               ACPM2918
 Assembly ID:  0x0b9c            Assembly Version:  01.10
 Date:         01-13-2017        Assembly Flags:    0x00
 Version:      REV 10            CLEI Code:         CMUCAH5CAA
```

```
  ID: QFX10016 FTC                FRU Model Number:  QFX10016-FAN-CTRL
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 9c 01 0a 52 45 56 20 31 30 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 33 30 39 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 4d 32 39 31 38 00 0d 01 07
   Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 35 43 41 41 51
   Address 0x50: 46 58 31 30 30 31 36 2d 46 41 4e 2d 43 54 52 4c
   Address 0x60: 00 00 00 00 00 00 41 41 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 6f ff ff ff ff ff ff ff ff ff ff ff ff
 FTC 1            REV 10   750-050309   ACPE8185          Fan Controller 16
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-050309        S/N:               ACPE8185
  Assembly ID:  0x0b9c            Assembly Version:  01.10
  Date:         12-22-2016        Assembly Flags:    0x00
  Version:      REV 10            CLEI Code:         CMUCAH5CAA
  ID: QFX10016 FTC                FRU Model Number:  QFX10016-FAN-CTRL
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 9c 01 0a 52 45 56 20 31 30 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 33 30 39 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 45 38 31 38 35 00 16 0c 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 35 43 41 41 51
   Address 0x50: 46 58 31 30 30 31 36 2d 46 41 4e 2d 43 54 52 4c
   Address 0x60: 00 00 00 00 00 00 41 41 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 6f ff ff ff ff ff ff ff ff ff ff ff ff
 Fan Tray 0       REV 10   760-077141   ACPV7288          Fan Tray 16
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          760-077141        S/N:               ACPV7288
  Assembly ID:  0x0bf1            Assembly Version:  01.10
  Date:         06-07-2017        Assembly Flags:    0x00
  Version:      REV 10            CLEI Code:         CMUCAH4CAA
  ID: QFX10016 FHB                FRU Model Number:  JNP10016-FAN
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b f1 01 0a 52 45 56 20 31 30 00 00
   Address 0x10: 00 00 00 00 37 36 30 2d 30 37 37 31 34 31 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 56 37 32 38 38 00 07 06 07
   Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 34 43 41 41 4a
   Address 0x50: 4e 50 31 30 30 31 36 2d 46 41 4e 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 0d ff ff ff ff ff ff ff ff ff ff ff ff
 Fan Tray 1       REV 10   760-057901   ACPL0546          Fan Tray 16
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          760-057901        S/N:               ACPL0546
  Assembly ID:  0x0bf1            Assembly Version:  01.10
  Date:         02-14-2017        Assembly Flags:    0x00
  Version:      REV 10            CLEI Code:         CMUCAH4CAA
  ID: QFX10016 FHB                FRU Model Number:  QFX10016-FAN
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b f1 01 0a 52 45 56 20 31 30 00 00
   Address 0x10: 00 00 00 00 37 36 30 2d 30 35 37 39 30 31 00 00
   Address 0x20: 53 2f 4e 20 41 43 50 4c 30 35 34 36 00 0e 02 07
```

```
  Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 34 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 46 41 4e 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 0d ff ff ff ff ff ff ff ff ff ff ff ff
SIB 0            REV 15   750-058270   ACPM2804         Switch Fabric 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-058270        S/N:               ACPM2804
 Assembly ID:  0x0bed            Assembly Version:  01.15
 Date:         12-21-2016        Assembly Flags:    0x00
 Version:      REV 15            CLEI Code:         CMUCAH6CAA
 ID: QFX10016 SIB                FRU Model Number:  QFX10016-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ed 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 38 32 37 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4d 32 38 30 34 00 15 0c 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 36 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 00 00 00 00 00 00 00 00 00 00 00 00
SIB 1            REV 15   750-058270   ACPM2808         Switch Fabric 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-058270        S/N:               ACPM2808
 Assembly ID:  0x0bed            Assembly Version:  01.15
 Date:         12-21-2016        Assembly Flags:    0x00
 Version:      REV 15            CLEI Code:         CMUCAH6CAA
 ID: QFX10016 SIB                FRU Model Number:  QFX10016-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ed 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 38 32 37 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4d 32 38 30 38 00 15 0c 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 36 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 00 00 00 00 00 00 00 00 00 00 00 00
SIB 2            REV 15   750-058270   ACPL4450         Switch Fabric 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-058270        S/N:               ACPL4450
 Assembly ID:  0x0bed            Assembly Version:  01.15
 Date:         02-17-2017        Assembly Flags:    0x00
 Version:      REV 15            CLEI Code:         CMUCAH6CAA
 ID: QFX10016 SIB                FRU Model Number:  QFX10016-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ed 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 38 32 37 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4c 34 34 35 30 00 11 02 07
  Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 36 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 00 00 00 00 00 00 00 00 00 00 00 00
SIB 3            REV 15   750-058270   ACPJ9834         Switch Fabric 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
```

```
 P/N:           750-058270        S/N:               ACPJ9834
 Assembly ID:  0x0bed            Assembly Version:  01.15
 Date:          12-17-2016        Assembly Flags:    0x00
 Version:       REV 15            CLEI Code:         CMUCAH6CAA
 ID: QFX10016 SIB                 FRU Model Number:  QFX10016-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ed 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 38 32 37 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4a 39 38 33 34 00 11 0c 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 36 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 00 00 00 00 00 00 00 00 00 00 00 00
SIB 4            REV 15   750-058270   ACPM2814          Switch Fabric 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:           750-058270        S/N:               ACPM2814
 Assembly ID:  0x0bed            Assembly Version:  01.15
 Date:          12-21-2016        Assembly Flags:    0x00
 Version:       REV 15            CLEI Code:         CMUCAH6CAA
 ID: QFX10016 SIB                 FRU Model Number:  QFX10016-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ed 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 38 32 37 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4d 32 38 31 34 00 15 0c 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 36 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 00 00 00 00 00 00 00 00 00 00 00 00
SIB 5            REV 15   750-058270   ACPL4277          Switch Fabric 16
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:           750-058270        S/N:               ACPL4277
 Assembly ID:  0x0bed            Assembly Version:  01.15
 Date:          02-17-2017        Assembly Flags:    0x00
 Version:       REV 15            CLEI Code:         CMUCAH6CAA
 ID: QFX10016 SIB                 FRU Model Number:  QFX10016-SF
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b ed 01 0f 52 45 56 20 31 35 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 38 32 37 30 00 00
  Address 0x20: 53 2f 4e 20 41 43 50 4c 34 32 37 37 00 11 02 07
  Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4d 55 43 41 48 36 43 41 41 51
  Address 0x50: 46 58 31 30 30 31 36 2d 53 46 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 42 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 00 00 00 00 00 00 00 00 00 00 00 00
FPD Board        REV 07   711-054687   ACPL1407          Front Panel Display
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:           711-054687        S/N:               ACPL1407
 Assembly ID:  0x0bf2            Assembly Version:  01.07
 Date:          02-12-2017        Assembly Flags:    0x00
 Version:       REV 07
 ID: QFX10000 FPD
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```



```
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 0b f2 01 07 52 45 56 20 30 37 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 35 34 36 38 37 00 00
 Address 0x20: 53 2f 4e 20 41 43 50 4c 31 34 30 37 00 0c 02 07
 Address 0x30: e1 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware models (MX104 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 20   750-044219   CAAS5849          PROTO-ASSEMBLY
PEM 0            REV 01   740-045932   1H072400065
Routing Engine 0 REV 16   750-044228   CAAR5915          PROTO-ASSEMBLY
AFEB 0                    BUILTIN      BUILTIN
FPC 0                     BUILTIN      BUILTIN
FPC 1                     BUILTIN      BUILTIN
  MIC 0          REV 01   750-046905   CAAK7103          MIC-3D-20GE-SFP-EH
FPC 2                     BUILTIN      BUILTIN
Fan Tray         REV 02   711-049570   CAAX6538          PROTO-ASSEMBLY
```

show chassis hardware models (PTX10008 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 27   750-054097   ACPD4307          QFX10008-CHAS
CB 0             REV 02   750-068820   ACNZ4440          QFX10000-RE
CB 1             REV 02   750-068820   ACNZ8284          QFX10000-RE
FPC 0            REV 36   750-051354   ACNP4679          QFX10000-36Q
  PIC 0                   BUILTIN      BUILTIN
FPC 1            REV 33   750-051354   ACNX8831          QFX10000-36Q
  PIC 0                   BUILTIN      BUILTIN
FPC 2            REV 32   750-051357   ACPB0341          QFX10000-30C
  PIC 0                   BUILTIN      BUILTIN
FPC 3            REV 35   750-051357   ACPD2186          QFX10000-30C
  PIC 0                   BUILTIN      BUILTIN
FPC 5            REV 08   750-068822   ACPF0057          QFX10000-36Q
  PIC 0                   BUILTIN      BUILTIN
FPC 6            REV 08   750-068822   ACPE9951          QFX10000-36Q
  PIC 0                   BUILTIN      BUILTIN
FPD Board        REV 07   711-054687   ACPC7142
Power Supply 0   REV 02   740-049388   1EDL62102N9       QFX10000-PWR-AC
Power Supply 1   REV 02   740-049388   1EDL60300KX       QFX10000-PWR-AC
Power Supply 2   REV 02   740-049388   1EDL60300DL       QFX10000-PWR-AC
Power Supply 3   REV 02   740-049388   1EDL61701BT       QFX10000-PWR-AC
Power Supply 4   REV 02   740-049388   1EDL62102P7       QFX10000-PWR-AC
Power Supply 5   REV 02   740-049388   1EDL62102PP       QFX10000-PWR-AC
FTC 0            REV 14   750-050108   ACPE4038          QFX10008-FAN-CTRL
FTC 1            REV 14   750-050108   ACPE4032          QFX10008-FAN-CTRL
Fan Tray 0       REV 09   760-054372   ACPD6799          QFX10008-FAN
Fan Tray 1       REV 09   760-054372   ACNZ3584          QFX10008-FAN
SIB 0            REV 24   750-050058   ACPD4587          QFX10008-SF
SIB 1            REV 24   750-050058   ACNZ0635          QFX10008-SF
SIB 2            REV 24   750-050058   ACPD4908          QFX10008-SF
SIB 3            REV 24   750-050058   ACNZ0617          QFX10008-SF
```

```
SIB 4            REV 24   750-050058   ACNZ0527         QFX10008-SF
SIB 5            REV 23   750-050058   ACNX6980         QFX10008-SF
```

show chassis hardware models (PTX10016 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number    FRU model number
Midplane         REV 24   750-077138   ACPR5157         JNP10016
CB 0             REV 04   711-065897   CAHA9983         PROTO-ASSEMBLY
CB 1             REV 05   711-065897   CAJD3802         PROTO-ASSEMBLY
FPC 2
  PIC 0                   BUILTIN      BUILTIN
FPC 4            REV 35   750-071976   ACPD2168         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 5            REV 13   750-068822   ACPA0336         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 6            REV 41   750-071976   ACPF0695         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 7            REV 35   750-071976   ACPD2139         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 8            REV 35   750-071976   ACPD2142         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 9            REV 41   750-071976   ACPM5461         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 10           REV 35   750-071976   ACNS6795         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 11           REV 35   750-071976   ACPD1831         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 13           REV 41   750-071976   ACPS2075         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
FPC 15           REV 37   750-071976   ACPL7163         JNP10K-LC1101
  PIC 0                   BUILTIN      BUILTIN
Power Supply 0   REV 01   740-073147   1EDM6171155      JNP10K-PWR-DC
Power Supply 1   REV 01   740-073147   1EDM6281575      JNP10K-PWR-DC
Power Supply 2   REV 01   740-073147   1EDM6171044      JNP10K-PWR-DC
Power Supply 3   REV 01   740-073147   1EDM6281244      JNP10K-PWR-DC
Power Supply 4   REV 01   740-073147   1EDM6282093      JNP10K-PWR-DC
Power Supply 5   REV 01   740-073147   1EDM6281413      JNP10K-PWR-DC
Power Supply 6   REV 01   740-073147   1EDM6171071      JNP10K-PWR-DC
Power Supply 7   REV 01   740-073147   1EDM6170709      JNP10K-PWR-DC
Power Supply 8   REV 01   740-073147   1EDM6171169      JNP10K-PWR-DC
Power Supply 9   REV 01   740-073147   1EDM6170754      JNP10K-PWR-DC
Fan Tray 0                                              QFX5100-FAN-AFO
Fan Tray 1                                              QFX5100-FAN-AFO
SIB 0            REV 15   750-077140   ACPV3933         JNP10016-SF
SIB 1            REV 15   750-077140   ACPV3938         JNP10016-SF
SIB 2            REV 15   750-077140   ACPV3974         JNP10016-SF
SIB 3            REV 15   750-077140   ACPV3879         JNP10016-SF
SIB 4            REV 15   750-077140   ACPV3964         JNP10016-SF
SIB 5            REV 15   750-077140   ACPV3981         JNP10016-SF
FPD Board        REV 07   711-054687   ACPS8855
```

show chassis hardware clei-models (MX104 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code        FRU model number
Midplane         REV 20   750-044219   PROTOXCLEI       PROTO-ASSEMBLY
PEM 0            REV 01   740-045932
```



```
Routing Engine 0 REV 16   750-044228   PROTOXCLEI       PROTO-ASSEMBLY
AFEB 0                    BUILTIN
FPC 0                     BUILTIN
FPC 1                     BUILTIN
  MIC 0          REV 01   750-046905   PROTOXCLEI       MIC-3D-20GE-SFP-EH
FPC 2                     BUILTIN
Fan Tray        REV 02   711-049570   CAAX6538         PROTO-ASSEMBLY
```

show chassis hardware (MX240 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN10C7F7EAFC     MX240
Midplane         REV 01   710-021041   TR1502           MX240 Backplane
FPM Board        REV 01   710-017254   KD4017           Front Panel Display
PEM 0            Rev 02   740-017330   000332           PS 1.2-1.7kW; 100-240V
AC in
PEM 1            Rev 02   740-017330   000226           PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 06   740-013063   1000703522       RE-S-2000
Routing Engine 1 REV 06   740-015113   1000687625       RE-S-1300
CB 0             REV 07   710-013385   KC9057           MX SCB
CB 1             REV 05   710-013385   JY4760           MX SCB
FPC 1            REV 01   750-021679   KC7340           DPCE 40x 1GE R
  CPU            REV 06   710-013713   KD4078           DPC PMB
  PIC 0                   BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 0       REV 01   740-011613   P9F18ME          SFP-SX
  PIC 1                   BUILTIN      BUILTIN          10x 1GE(LAN)
  PIC 2                   BUILTIN      BUILTIN          10x 1GE(LAN)
  PIC 3                   BUILTIN      BUILTIN          10x 1GE(LAN)
FPC 2            REV 04   710-016669   JS4529           DPCE 40x 1GE R EQ
  CPU            REV 06   710-013713   KB3969           DPC PMB
  PIC 0                   BUILTIN      BUILTIN          10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3Y79          SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3XU8          SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3YG6          SFP-SX
    Xcvr 3       REV 01   740-011613   PBG3XUG          SFP-SX
    Xcvr 4       REV 01   740-011613   PBG3XTJ          SFP-SX
  PIC 1                   BUILTIN      BUILTIN          10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3ZUM          SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3Y5H          SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3UZT          SFP-SX
    Xcvr 3       REV 01   740-011613   PBG3US1          SFP-SX
  PIC 2                   BUILTIN      BUILTIN          10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3YG7          SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3XZ9          SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3XTY          SFP-SX
    Xcvr 3       REV 01   740-011613   PBG3UZG          SFP-SX
  PIC 3                   BUILTIN      BUILTIN          10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3Y8W          SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3YVX          SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3YB3          SFP-SX
    Xcvr 3       REV 01   740-011613   PBG43VQ          SFP-SX
Fan Tray 0       REV 01   710-021113   JS4642           MX240 Fan Tray
```

show chassis hardware detail (MX 240 Router with Routing Engine Displaying DIMM Information)

```
user@host> show chassis hardware detail
```

```
Item             Version  Part number  Serial number     Description
Chassis                                JN11279B4AFC      MX240 Backplane
Midplane         REV 07   760-021404   TS2474            MX240 Backplane
FPM Board        REV 03   760-021392   XC2643            Front Panel Display
PEM 0            Rev 03   740-017343   QCS0908A068       DC Power Entry Module
Routing Engine 0 REV 01   740-031117   AARCH00           RE-S-1800x4
  ad0    3764 MB  STEC M2+ CF 9.0.2    STIM2Q3209239145303 Removable Compact Flash

  ad1   28626 MB  WDC SSD-F0030S-5000  C933Z036237215548S00 Compact Flash
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263E-F8S DIE REV-0 PCB REV-0      MFR ID-ce80
  DIMM 1         VL31B5263E-F8S DIE REV-0 PCB REV-0      MFR ID-ce80
  DIMM 2         VL31B5263E-F8S DIE REV-0 PCB REV-0      MFR ID-ce80
  DIMM 3         SL31B5263E-F8S DIE REV-0 PCB REV-0      MFR ID-ce80
CB 0             REV 03   710-021523   XD7225            MX SCB
Fan Tray 0       REV 01   710-021113   WZ4986            MX240 Fan Tray
```

show chassis hardware (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN10C7F7EAFC      MX240
Midplane         REV 01   710-021041   TR1502            MX240 Backplane
FPM Board        REV 01   710-017254   KD4017            Front Panel Display
PEM 0            Rev 02   740-017330   000332            PS 1.2-1.7kW; 100-240V
AC in
PEM 1            Rev 02   740-017330   000226            PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 06   740-013063   1000703522        RE-S-2000
Routing Engine 1 REV 06   740-015113   1000687625        RE-S-1300
CB 0             REV 02   710-031391   YE8494            Enhanced MX SCB
CB 1             REV 05   710-031391   YOP5764           Enhanced MX SCB
FPC 1            REV 01   750-021679   KC7340            DPCE 40x 1GE R
  CPU            REV 06   710-013713   KD4078            DPC PMB
  PIC 0                   BUILTIN      BUILTIN           10x 1GE(LAN)
    Xcvr 0       REV 01   740-011613   P9F18ME           SFP-SX
  PIC 1                   BUILTIN      BUILTIN           10x 1GE(LAN)
  PIC 2                   BUILTIN      BUILTIN           10x 1GE(LAN)
  PIC 3                   BUILTIN      BUILTIN           10x 1GE(LAN)
FPC 2            REV 04   710-016669   JS4529            DPCE 40x 1GE R EQ
  CPU            REV 06   710-013713   KB3969            DPC PMB
  PIC 0                   BUILTIN      BUILTIN           10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3Y79           SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3XU8           SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3YG6           SFP-SX
    Xcvr 3       REV 01   740-011613   PBG3XUG           SFP-SX
    Xcvr 4       REV 01   740-011613   PBG3XTJ           SFP-SX
  PIC 1                   BUILTIN      BUILTIN           10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3ZUM           SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3Y5H           SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3UZT           SFP-SX
    Xcvr 3       REV 01   740-011613   PBG3US1           SFP-SX
  PIC 2                   BUILTIN      BUILTIN           10x 1GE(LAN) EQ
    Xcvr 0       REV 01   740-011613   PBG3YG7           SFP-SX
    Xcvr 1       REV 01   740-011613   PBG3XZ9           SFP-SX
    Xcvr 2       REV 01   740-011613   PBG3XTY           SFP-SX
    Xcvr 3       REV 01   740-011613   PBG3UZG           SFP-SX
  PIC 3                   BUILTIN      BUILTIN           10x 1GE(LAN) EQ
```



```
    Xcvr 0        REV 01    740-011613   PBG3Y8W          SFP-SX
    Xcvr 1        REV 01    740-011613   PBG3YVX          SFP-SX
    Xcvr 2        REV 01    740-011613   PBG3YB3          SFP-SX
    Xcvr 3        REV 01    740-011613   PBG43VQ          SFP-SX
Fan Tray 0        REV 01    710-021113   JS4642           MX240 Fan Tray
```

show chassis hardware (MX480 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                 JN10C7F7FAFB     MX480
Midplane          REV 04   710-017414   TR2071           MX480 Midplane
FPM Board         REV 02   710-017254   KB8459           Front Panel Display
PEM 0             Rev 02   740-017330   QCS07519029      PS 1.2-1.7kW; 100-240V
AC in
PEM 1             Rev 02   740-017330   QCS07519041      PS 1.2-1.7kW; 100-240V
AC in
PEM 2             Rev 02   740-017330   QCS07519097      PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 07    740-013063   1000733381       RE-S-2000
Routing Engine 1 REV 07    740-013063   1000733540       RE-S-2000
CB 0              REV 07   710-013385   KA8022           MX SCB
CB 1              REV 07   710-013385   KA8303           MX SCB
FPC 0             REV 09   750-020452   KA8660           DPCE 40x 1GE X EQ
  CPU             REV 06   710-013713   KA8185           DPC PMB
  PIC 0                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
  PIC 1                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
  PIC 2                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
  PIC 3                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
Fan Tray                                                 Left Fan Tray
```

show chassis hardware (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                 JN10C7F7FAFB     MX480
Midplane          REV 04   710-017414   TR2071           MX480 Midplane
FPM Board         REV 02   710-017254   KB8459           Front Panel Display
PEM 0             Rev 02   740-017330   QCS07519029      PS 1.2-1.7kW; 100-240V
AC in
PEM 1             Rev 02   740-017330   QCS07519041      PS 1.2-1.7kW; 100-240V
AC in
PEM 2             Rev 02   740-017330   QCS07519097      PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 07    740-013063   1000733381       RE-S-2000
Routing Engine 1 REV 07    740-013063   1000733540       RE-S-2000
CB 0              REV 07   710-013385   KA8022           Enhanced MX SCB
CB 1              REV 07   710-013385   KA8303           Enhanced MX SCB
FPC 0             REV 09   750-020452   KA8660           DPCE 40x 1GE X EQ
  CPU             REV 06   710-013713   KA8185           DPC PMB
  PIC 0                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
  PIC 1                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
  PIC 2                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
  PIC 3                    BUILTIN      BUILTIN          10x 1GE(LAN) EQ
Fan Tray                                                 Left Fan Tray
```

show chassis hardware (MX480 Routers with MPC5E and Built-In OTN PIC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11C0338AFB      MX480
Midplane         REV 05   710-017414   ABAB8430          MX480 Midplane
FPM Board        REV 02   710-017254   ZS8005            Front Panel Display
PEM 0            Rev 05   740-029970   QCS1024U089       PS 1.4-2.52kW; 90-264V
AC in
PEM 1            Rev 10   740-029970   QCS1314U0FJ       PS 1.4-2.52kW; 90-264V
AC in
PEM 2            Rev 07   740-029970   QCS1121U076       PS 1.4-2.52kW; 90-264V
AC in
Routing Engine 0 REV 05   740-031116   9009092471        RE-S-1800x4
Routing Engine 1 REV 05   740-031116   9009097958        RE-S-1800x4
CB 0             REV 16   750-031391   CAAX0789          Enhanced MX SCB
CB 1             REV 16   750-031391   CAAX0856          Enhanced MX SCB
FPC 0            REV 32   750-028467   ABBP1782          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBP5410          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   983152A00038      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11F00211         SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AQ72LPB           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AHN0WR5           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   B11J03627         SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11F00300         SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQ42WSS           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AQ43HGC           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   ANA0ND0           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   ANA0NGF           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   ANA0NG9           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   ANA0MP9           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AQA06CG           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   19T511100493      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   APR040J           SFP+-10G-SR
FPC 1            REV 26   750-046005   CACN1894          MPC5E 3D Q 2CGE+4XGE
  CPU            REV 09   711-045719   CACN8698          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   163363A03046      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJ40JS8           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
  PIC 2                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   153363A00593      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJ40JUJ           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
    Xcvr 0                NON-JNPR     UQC0B53           CFP2-100G-LR4-D
FPC 2            REV 26   750-046005   CACN1891          MPC5E 3D Q 2CGE+4XGE
  CPU            REV 09   711-045719   CACN8694          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0                NON-JNPR     URA012A           SFP+-10G-LR
  PIC 1                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
    Xcvr 0                NON-JNPR     J13F47042         CFP2-100G-LR4-D
  PIC 2                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   AJC0BM3           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   11T511100917      SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
    Xcvr 0                NON-JNPR     UQK07SU           CFP2-100G-LR4-D
```



```
FPC 3            REV 03   750-045372   CAAD9425         MPCE Type 3 3D
CPU            REV 08   711-035209   CAAD9094         HMPC PMB 2G
  MIC 0          REV 14   750-033196   CAAW9204         1X100GE CXP
    PIC 0                 BUILTIN      BUILTIN          1X100GE CXP
      Xcvr 0     REV 01   740-046563   XD16FC034        CFP2-100G-SR10
  MIC 1          REV 19   750-033199   CAAJ1814         1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN          1X100GE CFP
FPC 4            REV 21.0.11 750-045715 CAAY3568        MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 07   711-045719   CAAW7430         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          12X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   AP406NG          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AR41NLP          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11D05630        SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          12X10GE SFPP OTN
  PIC 2                   BUILTIN      BUILTIN          3X40GE QSFPP
  PIC 3                   BUILTIN      BUILTIN          3X40GE QSFPP
  WAN MEZZ       REV 12   750-049136   CACM6678         MPC5E 24XGE OTN Mezz
FPC 5            REV 11   750-045372   CABK7539         MPCE Type 3 3D
  CPU            REV 08   711-035209   CABJ2466         HMPC PMB 2G
  MIC 0          REV 19   750-033199   CAAJ9719         1X100GE CFP
    PIC 0                 BUILTIN      BUILTIN          1X100GE CFP
      Xcvr 0     REV 01   740-035329   UP1020P          CFP-100G-SR10
  MIC 1          REV 07   750-033196   YZ0797           1X100GE CXP
    PIC 2                 BUILTIN      BUILTIN          1X100GE CXP
      Xcvr 0     REV 01   740-046563   XC42FC022        CFP2-100G-SR10
Fan Tray                                                Enhanced Left Fan Tray
```

show chassis hardware detail (MX480 Routers with MPC5E and Built-In OTN PIC)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN11C0338AFB     MX480
Midplane         REV 05   710-017414   ABAB8430         MX480 Midplane
FPM Board        REV 02   710-017254   ZS8005           Front Panel Display
PEM 0            Rev 05   740-029970   QCS1024U089      PS 1.4-2.52kW; 90-264V
AC in
PEM 1            Rev 10   740-029970   QCS1314U0FJ      PS 1.4-2.52kW; 90-264V
AC in
PEM 2            Rev 07   740-029970   QCS1121U076      PS 1.4-2.52kW; 90-264V
AC in
Routing Engine 0 REV 05   740-031116   9009092471       RE-S-1800x4
  ad0     3896 MB  VRFCF14096DIHK1     VM4096MB 6862    Compact Flash
  ad1    30533 MB  UGB94ARF32H0S3-KC   UNIGEN-478612-001127 Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel            uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087    uhub1
  DIMM 0         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 05   740-031116   9009097958       RE-S-1800x4
  ad0     3896 MB  VRFCF14096DIHK1     VM4096MB 6145    Compact Flash
  ad1    30533 MB  UGB94ARF32H0S3-KC   UNIGEN-499551-000273 Disk 1
CB 0             REV 16   750-031391   CAAX0789         Enhanced MX SCB
CB 1             REV 16   750-031391   CAAX0856         Enhanced MX SCB
FPC 0            REV 32   750-028467   ABBP1782         MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBP5410         AMPC PMB
  PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   983152A00038     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11F00211        SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AQ72LPB          SFP+-10G-SR
```

```
    Xcvr 3       REV 01   740-031980   AHNOWR5          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   B11J03627        SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11F00300        SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQ42WSS          SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AQ43HGC          SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   ANA0ND0          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   ANA0NGF          SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   ANA0NG9          SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   ANA0MP9          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AQA06CG          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   19T511100493     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   APR040J          SFP+-10G-SR
FPC 1            REV 26   750-046005   CACN1894         MPC5E 3D Q 2CGE+4XGE
  CPU            REV 09   711-045719   CACN8698         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   163363A03046     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJ40JS8          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
  PIC 2                   BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   153363A00593     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJ40JUJ          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0                NON-JNPR     UQC0B53          CFP2-100G-LR4-D
FPC 2            REV 26   750-046005   CACN1891         MPC5E 3D Q 2CGE+4XGE
  CPU            REV 09   711-045719   CACN8694         RMPC PMB
PIC 0                   BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0                NON-JNPR     URA012A          SFP+-10G-LR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0                NON-JNPR     J13F47042        CFP2-100G-LR4-D
  PIC 2                   BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   AJC0BM3          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   11T511100917     SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0                NON-JNPR     UQK07SU          CFP2-100G-LR4-D
FPC 3            REV 03   750-045372   CAAD9425         MPCE Type 3 3D
  CPU            REV 08   711-035209   CAAD9094         HMPC PMB 2G
  MIC 0          REV 14   750-033196   CAAW9204         1X100GE CXP
    PIC 0                 BUILTIN      BUILTIN          1X100GE CXP
      Xcvr 0     REV 01   740-046563   XD16FC034        CFP2-100G-SR10
  MIC 1          REV 19   750-033199   CAAJ1814         1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN          1X100GE CFP
FPC 4            REV 21.0.11 750-045715 CAAY3568        MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 07   711-045719   CAAW7430         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          12X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   AP406NG          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AR41NLP          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11D05630        SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          12X10GE SFPP OTN
  PIC 2                   BUILTIN      BUILTIN          3X40GE QSFPP
  PIC 3                   BUILTIN      BUILTIN          3X40GE QSFPP
  WAN MEZZ       REV 12   750-049136   CACM6678         MPC5E 24XGE OTN Mezz
FPC 5            REV 11   750-045372   CABK7539         MPCE Type 3 3D
  CPU            REV 08   711-035209   CABJ2466         HMPC PMB 2G
  MIC 0          REV 19   750-033199   CAAJ9719         1X100GE CFP
    PIC 0                 BUILTIN      BUILTIN          1X100GE CFP
      Xcvr 0     REV 01   740-035329   UP1020P          CFP-100G-SR10
  MIC 1          REV 07   750-033196   YZ0797           1X100GE CXP
    PIC 2                 BUILTIN      BUILTIN          1X100GE CXP
```



```
      Xcvr 0      REV 01    740-046563    XC42FC022         CFP2-100G-SR10
Fan Tray                                                    Enhanced Left Fan Tray
```

show chassis hardware extensive (MX480 Routers with MPC5E and Built-In OTN PIC)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11C0338AFB      MX480
 Jedec Code:   0x7fb0             EEPROM Version:    0x02
                                  S/N:               JN11C0338AFB
 Assembly ID:  0x01fe             Assembly Version:  00.00
 Date:         00-00-0000         Assembly Flags:    0x02
 ID: MX480
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 01 fe 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 31 43 30 33 33 38 41 46 42 02 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 05   710-017414   ABAB8430          MX480 Midplane
 Jedec Code:   0x7fb0             EEPROM Version:    0x01
 P/N:          710-017414         S/N:               ABAB8430
 Assembly ID:  0x01fe             Assembly Version:  01.05
 Date:         12-13-2011         Assembly Flags:    0x00
 Version:      REV 05
 ID: MX480 Midplane               FRU Model Number:  CHAS-BP-MX480-S
 Board Information Record:
  Address 0x00: ad 01 08 00 00 23 9c fc 98 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 01 fe 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 34 31 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 38 34 33 30 00 0d 0c 07
  Address 0x30: db ff ff ff ad 01 08 00 00 23 9c fc 98 00 ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
  Address 0x50: 48 41 53 2d 42 50 2d 4d 58 34 38 30 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 02   710-017254   ZS8005            Front Panel Display
 Jedec Code:   0x7fb0             EEPROM Version:    0x01
 P/N:          710-017254         S/N:               ZS8005
 Assembly ID:  0x01ff             Assembly Version:  01.02
 Date:         11-21-2011         Assembly Flags:    0x00
 Version:      REV 02
 ID: Front Panel Display          FRU Model Number:  CRAFT-MX480-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 01 ff 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 32 35 34 00 00
  Address 0x20: 53 2f 4e 20 5a 53 38 30 30 35 00 00 00 15 0b 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
  Address 0x50: 52 41 46 54 2d 4d 58 34 38 30 2d 53 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
```



```
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PEM 0            Rev 05   740-029970   QCS1024U089       PS 1.4-2.52kW; 90-264V
AC in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-029970        S/N:               QCS1024U089
 Assembly ID:  0x0432            Assembly Version:  01.05
 Date:         06-17-2010        Assembly Flags:    0x00
 Version:      Rev 05
 ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number:  PWR-MX480-2520-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 32 01 05 52 65 76 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
  Address 0x20: 51 43 53 31 30 32 34 55 30 38 39 00 00 11 06 07
  Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1            Rev 10   740-029970   QCS1314U0FJ       PS 1.4-2.52kW; 90-264V
AC in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-029970        S/N:               QCS1314U0FJ
 Assembly ID:  0x0432            Assembly Version:  01.10
 Date:         04-04-2013        Assembly Flags:    0x00
 Version:      Rev 10
 ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number:  PWR-MX480-2520-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 32 01 0a 52 65 76 20 31 30 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
  Address 0x20: 51 43 53 31 33 31 34 55 30 46 4a 00 00 04 04 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 2            Rev 07   740-029970   QCS1121U076       PS 1.4-2.52kW; 90-264V
AC in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-029970        S/N:               QCS1121U076
 Assembly ID:  0x0432            Assembly Version:  01.07
 Date:         05-23-2011        Assembly Flags:    0x00
 Version:      Rev 07
 ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number:  PWR-MX480-2520-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 32 01 07 52 65 76 20 30 37 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
  Address 0x20: 51 43 53 31 31 32 31 55 30 37 36 00 00 17 05 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 05   740-031116   9009092471        RE-S-1800x4
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-031116        S/N:               9009092471
```



```
 Assembly ID:  0x09c0            Assembly Version:  01.05
 Date:         11-01-2011        Assembly Flags:    0x00
 Version:      REV 05            CLEI Code:         COUCALDBAA
 ID: RE-S-1800x4                 FRU Model Number:  RE-S-1800X4-16G-S
 Board Information Record:
  Address 0x00: 54 32 30 32 37 43 41 2d 34 32 46 42 23 23 23 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 30 39 32 34 37 31 00 00 00 01 0b 07
  Address 0x30: db ff ff ff 54 32 30 32 37 43 41 2d 34 32 46 42
  Address 0x40: 23 23 23 00 01 43 4f 55 43 41 4c 44 42 41 41 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4b ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3896 MB  VRFCF14096DIHK1      VM4096MB 6862     Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-478612-001127 Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 05   740-031116   9009097958        RE-S-1800x4
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-031116        S/N:               9009097958
 Assembly ID:  0x09c0            Assembly Version:  01.05
 Date:         02-06-2012        Assembly Flags:    0x00
 Version:      REV 05            CLEI Code:         COUCALDBAA
 ID: RE-S-1800x4                 FRU Model Number:  RE-S-1800X4-16G-S
 Board Information Record:
  Address 0x00: 54 32 30 32 37 43 41 2d 34 32 46 42 23 23 23 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 30 39 37 39 35 38 00 00 00 06 02 07
  Address 0x30: dc ff ff ff 54 32 30 32 37 43 41 2d 34 32 46 42
  Address 0x40: 23 23 23 00 01 43 4f 55 43 41 4c 44 42 41 41 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4b ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3896 MB  VRFCF14096DIHK1      VM4096MB 6145     Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-499551-000273 Disk 1

...
```

show chassis hardware (MX960 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                                  MX960
Midplane         REV 01   710-013698   AA6082            MX960 Midplane
PIM              Rev 01   740-013110   000008            Power Inlet Module
PEM 2
PEM 3            Rev 01   740-013682   000038            PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 00   740-015113   1000617944        RE-S-1300
CB 0             REV 05   710-013725   JK6947            MX960 Test SCB
FPC 4            REV 01   710-013305   JM7617            MX960 Test DPC
  CPU
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN)
```

```
  PIC 1                        BUILTIN      BUILTIN            10x 1GE
FPC 7             REV 01       710-013305   JL9634             MX960 Test DPC
  CPU
  PIC 0                        BUILTIN      BUILTIN            1x 10GE(LAN/WAN)
    Xcvr 0                     NON-JNPR     MYBG65I82C         XFP-10G-SR
  PIC 1                        BUILTIN      BUILTIN            10x 1GE
    Xcvr 1        REV 01       740-011782   P7N0368            SFP-SX
    Xcvr 4        REV 01       740-011782   P8J1W27            SFP-SX
    Xcvr 6        REV 01       740-011782   P8J1VSD            SFP-SX
    Xcvr 9        REV 01       740-011782   P8J1W25            SFP-SX
Fan Tray 0
Fan Tray 1
```

show chassis hardware (MX960 Router with Bidirectional Optics)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number      Description
Chassis                                 JN10BA5B9AFA       MX960
Midplane          REV 03   710-013698   TR0234             MX960 Backplane
FPM Board         REV 03   710-014974   JA0878             Front Panel Display
PDM              Rev 03   740-013110   QCS11135028        Power Distribution Module
PEM 0             Rev 03   740-013682   QCS11154036        PS 1.7kW; 200-240VAC in
PEM 1             Rev 03   740-013682   QCS11154010        PS 1.7kW; 200-240VAC in
PEM 2             Rev 03   740-013682   QCS11154022        PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 06   740-013063   1000691458         RE-S-2000
CB 0              REV 07   710-013385   KA2190             MX SCB
CB 1              REV 07   710-013385   KA0837             MX SCB
FPC 3             REV 02   750-018122   KB3890             DPCE 40x 1GE R
  CPU
FPC 4             REV 01   750-018122   KB3889             DPCE 40x 1GE R
  CPU             REV 06   710-013713   KB3976             DPC PMB
  PIC 0                    BUILTIN      BUILTIN            10x 1GE(LAN)
    Xcvr 1        REV 01   740-020426   4910549            SFP-1000BASE-BX40-D
    Xcvr 2        REV 01   740-020426   4910551            SFP-1000BASE-BX40-D
    Xcvr 5        REV 01   740-021340   77E245N00006       SFP-1000BASE-BX10-U
    Xcvr 6        REV 01   740-020425   4882821            SFP-1000BASE-BX40-U
    Xcvr 8        REV 01   740-020425   4882820            SFP-1000BASE-BX40-U
  PIC 1                    BUILTIN      BUILTIN            10x 1GE(LAN)
    Xcvr 0        REV 01   740-020465   77E555N00894       SFP-1000BASE-BX10-D
    Xcvr 1        REV 01   740-020465   75E467X00818       SFP-1000BASE-BX10-D
    Xcvr 2        REV 01   740-020465   75E467X00573       SFP-1000BASE-BX10-D
    Xcvr 3        REV 01   740-020465   4888227            SFP-1000BASE-BX10-D
    Xcvr 4        REV 01   740-020465   4888241            SFP-1000BASE-BX10-D
    Xcvr 5        REV 01   740-021340   77E245N00005       SFP-1000BASE-BX10-U
    Xcvr 6        REV 01   740-021340   76E245X00487       SFP-1000BASE-BX10-U
    Xcvr 7        REV 01   740-021341   5255889            SFP-1000BASE-BX10-U
    Xcvr 8        REV 01   740-021341   5255887            SFP-1000BASE-BX10-U
    Xcvr 9        REV 01   740-021340   77E245N00004       SFP-1000BASE-BX10-U
  PIC 2                    BUILTIN      BUILTIN            10x 1GE(LAN)
    Xcvr 0        REV 01   740-020424   5007582            SFP-1000BASE-BX10-D
    Xcvr 1        REV 01   740-020424   4888187            SFP-1000BASE-BX10-D
    Xcvr 2        REV 01   740-020424   4656500            SFP-1000BASE-BX10-D
    Xcvr 5        REV 01   740-021341   5255886            SFP-1000BASE-BX10-U
    Xcvr 7        REV 01   740-021340   77E245N00003       SFP-1000BASE-BX10-U
    Xcvr 8        REV 01   740-021341   5255888            SFP-1000BASE-BX10-U
  PIC 3                    BUILTIN      BUILTIN            10x 1GE(LAN)
    Xcvr 0        REV 01   740-017726   74S184H30341       SFP-EX
    Xcvr 1        REV 01   740-017726   4814061            SFP-EX
    Xcvr 5        REV 01   740-017726   6ZS184H31108       SFP-EX
    Xcvr 9        REV 01   740-021340   76E245X00486       SFP-1000BASE-BX10-U
```



```
Fan Tray 0
Fan Tray 1        REV 03   740-014971   TP0850           Fan Tray
```

show chassis hardware (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                 JN1096805AFA     MX960
Midplane          REV 03   710-013698   TR0183           MX960 Backplane
Fan Extender      REV 02   710-018051   JY5227           Extended Cable Manager
FPM Board         REV 03   710-014974   JZ6876           Front Panel Display
PDM              Rev 03   740-013110   QCS11035023      Power Distribution Module
PEM 1             Rev 03   740-013682   QCS1109400L      PS 1.7kW; 200-240VAC in
PEM 2             Rev 03   740-013682   QCS11094015      PS 1.7kW; 200-240VAC in
PEM 3             Rev 03   740-013682   QCS11094012      PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 06   740-013063   1000687969       RE-S-2000
Routing Engine 1 REV 06   740-013063   1000687955       RE-S-2000
CB 0              REV 11   750-031391   YZ6072           Enhanced MX SCB
CB 1              REV 11   750-031391   YZ6068           Enhanced MX SCB
CB 2              REV 11   750-031391   YZ6081           Enhanced MX SCB
FPC 0             REV 01   750-018122   KA5576           DPCE 40x 1GE R
  CPU             REV 06   710-013713   KB3961           DPC PMB
  PIC 0                    BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 0        REV 01   740-011613   P9F18GF          SFP-SX
    Xcvr 2        REV 01   740-011782   P9M0TL9          SFP-SX
    Xcvr 7        REV 01   740-011782   P9P0XXH          SFP-SX
    Xcvr 9        REV 01   740-011782   P9M0TN1          SFP-SX
  PIC 1                    BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 0        REV 01   740-011613   PAJ4UHC          SFP-SX
  PIC 2                    BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 0        REV 01   740-011613   PFF2CD0          SFP-SX
    Xcvr 1        REV 01   740-011613   PBG3ZUT          SFP-SX
    Xcvr 2        REV 01   740-011613   PFF2DDV          SFP-SX
    Xcvr 5        REV 01   740-011613   P8E2SST          SFP-SX
    Xcvr 9        REV 01   740-011782   PB8329N          SFP-SX
  PIC 3                    BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 0        REV 01   740-026192   1U0201084503342  SFP-100BASE-BX10-U
    Xcvr 1        REV 01   740-026193   1U1201084503313  SFP-100BASE-BX10-D
    Xcvr 2        REV 01   740-011613   PAJ4Y5B          SFP-SX
    Xcvr 6        REV 01   740-011782   P9M0U3M          SFP-SX
    Xcvr 7        REV 01   740-011782   P9M0TLA          SFP-SX
FPC 1             REV 16   750-031089   YL0719           MPC Type 2 3D
  CPU             REV 06   711-030884   YL1463           MPC PMB 2G
  MIC 0           REV 07   750-028387   JR6500           3D 4x 10GE  XFP
    PIC 0                  BUILTIN      BUILTIN          2x 10GE  XFP
      Xcvr 0      REV 01   740-014279   733019A00154     XFP-10G-LR
      Xcvr 1      REV 02   740-014289   T09F55034        XFP-10G-SR
    PIC 1                  BUILTIN      BUILTIN          2x 10GE  XFP
      Xcvr 0      REV 01   740-014279   913019B00791     XFP-10G-LR
      Xcvr 1      REV 01   740-014289   98S803A90384     XFP-10G-SR
  MIC 1           REV 24   750-028387   YJ3950           3D 4x 10GE  XFP
    PIC 2                  BUILTIN      BUILTIN          2x 10GE  XFP
      Xcvr 0      REV 02   740-014279   T10B36134        XFP-10G-LR
      Xcvr 1      REV 01   740-014289   T07M86354        XFP-10G-SR
    PIC 3                  BUILTIN      BUILTIN          2x 10GE  XFP
FPC 2             REV 08   710-014219   JY9654           DPCE 4x 10GE R
  CPU             REV 06   710-013713   JZ6549           DPC PMB
  PIC 0                    BUILTIN      BUILTIN          1x 10GE(LAN/WAN)
  PIC 1                    BUILTIN      BUILTIN          1x 10GE(LAN/WAN)
  PIC 2                    BUILTIN      BUILTIN          1x 10GE(LAN/WAN)
```

```
    Xcvr 0        REV 03   740-011571   C931BK028        XFP-10G-SR
  PIC 3                    BUILTIN      BUILTIN          1x 10GE(LAN/WAN)
FPC 3             REV 10   750-024199   XJ6692           MX FPC Type 3
  CPU             REV 03   710-022351   XF5182           DPC PMB
  PIC 0           REV 17   750-009553   RJ2945           4x OC-48 SONET
    Xcvr 1        REV 01   740-011785   PCP3YLL          SFP-SR
    Xcvr 3        REV 01   740-011785   PDS0MRY          SFP-SR
  PIC 1           REV 32   750-003700   DP2113           1x OC-192 12xMM VSR
FPC 5             REV 25   750-028467   YM8256           MPC 3D 16x 10GE
  CPU             REV 10   711-029089   YL3029           AMPC PMB
  PIC 0                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 1        REV 01   740-031980   AHN0X1Z          SFP+-10G-SR
  PIC 1                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
  PIC 2                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
  PIC 3                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
FPC 7             REV 02   750-031092   JR6658           MPC Type 1 3D Q
  CPU             REV 01   711-030884   JZ9038           MPC PMB 2G
  MIC 0           REV 08   750-028392   JZ8737           3D 20x 1GE(LAN) SFP
    PIC 0                  BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0      REV 01   740-011782   PBE2C6Y          SFP-SX
      Xcvr 2               NON-JNPR     U8105N8          SFP-SX
      Xcvr 4      REV 01   740-011613   PFM18EF          SFP-SX
      Xcvr 7      REV 01   740-011613   PFF2AM8          SFP-SX
      Xcvr 8      REV 01   740-011613   PFF2CT6          SFP-SX
    PIC 1                  BUILTIN      BUILTIN          10x 1GE(LAN) SFP
      Xcvr 0      REV 01   740-011782   PB82VHH          SFP-SX
      Xcvr 1      REV 01   740-011613   PFF2CSW          SFP-SX
      Xcvr 9      REV 01   740-011613   PFF2BY0          SFP-SX
  QXM 0           REV 04   711-028408   JR6372           MPC QXM
FPC 8             REV 05   750-024387   JW9754           MX FPC Type 2
  CPU             REV 03   710-022351   KF1651           DPC PMB
  PIC 0           REV 08   750-014730   DM3664           4x OC-3 1x OC-12 SFP
    Xcvr 0        REV 01   740-016065   81S290N00077     SFP-SR
    Xcvr 1                 NON-JNPR     2191844          SFP-SR
    Xcvr 2        REV 01   740-011618   PD81EE5          SFP-IR
  PIC 1           REV 08   750-014637   DM3671           4x OC-12-3 SFP
    Xcvr 0        REV 01   740-011785   PCK3UNK          SFP-SR
    Xcvr 3        REV 01   740-011785   PDS0MPZ          SFP-SR
FPC 10            REV 04   710-013699   JY4654           DPCE 40x 1GE R
  CPU             REV 05   710-013713   JS9717           DPC PMB
  PIC 0                    BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 5        REV 01   740-011782   PAR1L72          SFP-SX
    Xcvr 6        REV 01   740-011782   P8N1YQ4          SFP-SX
  PIC 1                    BUILTIN      BUILTIN          10x 1GE(LAN)
  PIC 2                    BUILTIN      BUILTIN          10x 1GE(LAN)
    Xcvr 0        REV 01   740-011782   P8Q2AVL          SFP-SX
    Xcvr 5        REV 01   740-011782   PAR1L7B          SFP-SX
    Xcvr 6        REV 01   740-011782   PAR1L2J          SFP-SX
    Xcvr 8        REV 01   740-011782   P8N1YMY          SFP-SX
  PIC 3                    BUILTIN      BUILTIN          10x 1GE(LAN)
Fan Tray 0        REV 03   740-014971   TP0567           Fan Tray
Fan Tray 1        REV 03   740-014971   TP0702           Fan Tray
```

show chassis hardware models (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware models
Hardware inventory:
Item              Version  Part number  Serial number    FRU model number
Midplane          REV 03   710-013698   TR0183           CHAS-BP-MX960-S
Fan Extender      REV 02   710-018051   JY5227           ECM-MX960
FPM Board         REV 03   710-014974   JZ6876           CRAFT-MX960-S
```

```
Routing Engine 0 REV 06   740-013063   1000687969       RE-S-2000-4096-S
Routing Engine 1 REV 06   740-013063   1000687955       RE-S-2000-4096-S
CB 0             REV 11   750-031391   YZ6072           SCBE-MX-S
CB 1             REV 11   750-031391   YZ6068           SCBE-MX-S
CB 2             REV 11   750-031391   YZ6081           SCBE-MX-S
FPC 0            REV 01   750-018122   KA5576           DPCE-R-40GE-SFP
FPC 1            REV 16   750-031089   YL0719           MX-MPC2-3D
  MIC 0          REV 07   750-028387   JR6500           MIC-3D-4XGE-XFP
  MIC 1          REV 24   750-028387   YJ3950           MIC-3D-4XGE-XFP
FPC 2            REV 08   710-014219   JY9654           DPC-R-4XGE-XFP
FPC 3            REV 10   750-024199   XJ6692           MX-FPC3
  PIC 0          REV 17   750-009553   RJ2945           PC-4OC48-SON-SFP
  PIC 1          REV 32   750-003700   DP2113           PC-1OC192-SON-VSR
FPC 5            REV 25   750-028467   YM8256           MPC-3D-16XGE-SFPP
FPC 7            REV 02   750-031092   JR6658           MX-MPC1-3D-Q
  MIC 0          REV 08   750-028392   JZ8737           MIC-3D-20GE-SFP
FPC 8            REV 05   750-024387   JW9754           MX-FPC2
  PIC 0          REV 08   750-014730   DM3664           PB-4OC3-1OC12-SON2-SFP
  PIC 1          REV 08   750-014637   DM3671           PB-4OC3-4OC12-SON-SFP
FPC 10           REV 04   710-013699   JY4654           DPC-R-40GE-SFP
Fan Tray 0       REV 03   740-014971   TP0567           FFANTRAY-MX960-S
Fan Tray 1       REV 03   740-014971   TP0702           FFANTRAY-MX960-S
```

show chassis hardware (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN1214852AFA     MX960
Midplane         REV 01   710-030012   ACAX3674         MX960 Backplane
FPM Board        REV 03   710-014974   CAAZ9326         Front Panel Display
PDM             Rev 03   740-013110   QCS17025017      Power Distribution Module
PEM 0            Rev 10   740-027760   QCS1702N062      PS 4.1kW; 200-240V AC
in
PEM 1            Rev 04   740-027760   QCS1422N02C      PS 4.1kW; 200-240V AC
in
PEM 2            Rev 09   740-027760   QCS1614N01X      PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 08   740-031116   9009131803       RE-S-1800x4
Routing Engine 1 REV 08   740-031116   9009124913       RE-S-1800x4
CB 0             REV 18   750-031391   CABF0579         Enhanced MX SCB
CB 1             REV 16   750-031391   CAAZ2471         Enhanced MX SCB
CB 2             REV 16   750-031391   CAAW9595         Enhanced MX SCB
FPC 0            REV 18   750-046005   CACE6574         MPC5E 3D Q 2CGE+4XGE
  CPU            REV 09   711-045719   CACG8908         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0       REV 01   740-021308   AQA0DYT          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQG0MS7          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0       REV 01   740-046563   XD16FC03Z        CFP2-100G-SR10
  PIC 2                   BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0       REV 01   740-021308   ANA0NAJ          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQG0MRQ          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0       REV 01   740-049775   J13K72993        CFP2-100G-LR4
FPC 1            REV 11   750-045372   CABK8154         MPCE Type 3 3D
  CPU            REV 08   711-035209   CABE7370         HMPC PMB 2G
  MIC 0          REV 07   750-033307   CABD5255         10X10GE SFPP
    PIC 0                 BUILTIN      BUILTIN          10X10GE SFPP
      Xcvr 0     REV 01   740-021308   AQ50319          SFP+-10G-SR
      Xcvr 1     REV 01   740-021308   AQ5035V          SFP+-10G-SR
```



```
        Xcvr 2      REV 01   740-021308   AQ502XJ          SFP+-10G-SR
        Xcvr 3      REV 01   740-021308   AQ43HHR          SFP+-10G-SR
        Xcvr 4      REV 01   740-021308   AQ502YA          SFP+-10G-SR
        Xcvr 5      REV 01   740-021308   AQ502EU          SFP+-10G-SR
        Xcvr 6      REV 01   740-021308   AQ502HR          SFP+-10G-SR
        Xcvr 7      REV 01   740-021308   AQ502A6          SFP+-10G-SR
        Xcvr 8      REV 01   740-021308   AQ43H8M          SFP+-10G-SR
    MIC 1           REV 14   750-033196   CAAP1398         1X100GE CXP
      PIC 2                  BUILTIN      BUILTIN          1X100GE CXP
        Xcvr 0      REV 01   740-046563   XD16FC064        CFP-100G-SR10
  FPC 3             REV 35   750-028467   CAAT9156         MPC 3D 16x 10GE
    CPU             REV 11   711-029089   CAAV4645         AMPC PMB
    PIC 0                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-021308   AQ43HZ1          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQ43HZC          SFP+-10G-SR
      Xcvr 2        REV 01   740-021308   AQ43HD2          SFP+-10G-SR
      Xcvr 3        REV 01   740-021308   AQ502HN          SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-021308   AQ43HGF          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQ501RZ          SFP+-10G-SR
      Xcvr 2        REV 01   740-021308   AQ5029V          SFP+-10G-SR
      Xcvr 3        REV 01   740-021308   AQ501X9          SFP+-10G-SR
    PIC 2                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-021308   AQ502ZN          SFP+-10G-SR
   Xcvr 1        REV 01   740-021308   AQ43H86          SFP+-10G-SR
      Xcvr 2        REV 01   740-021308   AQ502ZY          SFP+-10G-SR
      Xcvr 3        REV 01   740-021308   AQ502PZ          SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-021308   AQ503E6          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQ502XN          SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   B11F00213        SFP+-10G-SR
      Xcvr 3        REV 01   740-021308   AQ50336          SFP+-10G-SR
  FPC 4             REV 18   750-046005   CACE6568         MPC5E 3D Q 2CGE+4XGE
    CPU             REV 09   711-045719   CACG8900         RMPC PMB
    PIC 0                    BUILTIN      BUILTIN          2X10GE SFPP OTN
      Xcvr 0        REV 01   740-021308   AQA095A          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQG0M1E          SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN          1X100GE CFP2 OTN
      Xcvr 0                 NON-JNPR     FE13F000F        CFP2-100G-SR10
    PIC 2                    BUILTIN      BUILTIN          2X10GE SFPP OTN
      Xcvr 0        REV 01   740-021308   AQG0LYC          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQG0LYB          SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN          1X100GE CFP2 OTN
      Xcvr 0        REV 01   740-048813   XD32FE00Z        CFP2-100G-SR10
  FPC 5             REV 18   750-046005   CACE6577         MPC5E 3D Q 2CGE+4XGE
    CPU             REV 09   711-045719   CACG8902         RMPC PMB
    PIC 0                    BUILTIN      BUILTIN          2X10GE SFPP OTN
      Xcvr 0        REV 01   740-021308   AQG0MXE          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQG0LVY          SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN          1X100GE CFP2 OTN
      Xcvr 0        REV 01   740-046563   XD16FC03T        CFP2-100G-SR10
    PIC 2                    BUILTIN      BUILTIN          2X10GE SFPP OTN
      Xcvr 0        REV 01   740-021308   AQG0LW1          SFP+-10G-SR
      Xcvr 1        REV 01   740-021308   AQG0LW3          SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN          1X100GE CFP2 OTN
      Xcvr 0                 NON-JNPR     FE13F000J        CFP2-100G-SR10
  FPC 7             REV 09   750-037355   CAAF0937         MPC4E 3D 2CGE+8XGE
    CPU             REV 08   711-035209   CAAD8004         HMPC PMB 2G
    PIC 0                    BUILTIN      BUILTIN          4x10GE SFPP
      Xcvr 0        REV 01   740-021308   ANA0MM3          SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN          1X100GE CFP
```

```
    Xcvr 0       REV 01   740-035329   X000C163          CFP-100G-SR10
  PIC 2                   BUILTIN      BUILTIN           4x10GE SFPP
    Xcvr 0       REV 01   740-021308   AQG0MS6           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQG0MRX           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQG0M6Y           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AQG0LZM           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1X100GE CFP
    Xcvr 0       REV 01   740-035329   X12J00499         CFP-100G-SR10
FPC 8            REV 39   750-045715   CACD1903          MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 09   711-045719   CACD1815          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           12X10GE SFPP OTN
  PIC 1                   BUILTIN      BUILTIN           12X10GE SFPP OTN
  PIC 2                   BUILTIN      BUILTIN           3X40GE QSFPP
    Xcvr 0       REV 01   740-046565   QC480289          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-046565   QC480274          QSFP+-40G-SR4
    Xcvr 2       REV 01   740-046565   QD130190          QSFP+-40G-SR4
  PIC 3                   BUILTIN      BUILTIN           3X40GE QSFPP
    Xcvr 0       REV 01   740-046565   QD130197          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-046565   QD130180          QSFP+-40G-SR4
 Xcvr 2       REV 01   740-046565   QD130199          QSFP+-40G-SR4
  WAN MEZZ       REV 09   750-049136   CABN0415          MPC5E 24XGE OTN Mezz
FPC 9            REV 05   750-044444   CAAY9801          MPCE Type 2 3D P
  CPU            REV 04   711-038484   CAAW3673          MPCE PMB 2G
  MIC 0          REV 28   750-028387   CAAX1071          3D 4x 10GE  XFP
    PIC 0                 BUILTIN      BUILTIN           2x 10GE  XFP
      Xcvr 0              NON-JNPR     T12L92342         XFP-10G-SR
      Xcvr 1              NON-JNPR     T12L92303         XFP-10G-SR
    PIC 1                 BUILTIN      BUILTIN           2x 10GE  XFP
      Xcvr 0              NON-JNPR     CC07BK02X         XFP-10G-SR
  QXM 0          REV 06   711-028408   CAAW4883          MPC QXM
  QXM 1          REV 06   711-028408   CAAW4603          MPC QXM
FPC 10           REV 21.0.11 750-045715 CAAY3541         MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 07   711-045719   CAAW7426          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           12X10GE SFPP
    Xcvr 0       REV 01   740-031980   AHK01AP           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQ502ZU           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AP41BLS           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AQA08YA           SFP+-10G-SR
    Xcvr 4       REV 01   740-021308   AQA0K26           SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AQA06S3           SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   AQA06AS           SFP+-10G-SR
    Xcvr 8       REV 01   740-021308   AQA053N           SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   AQA0E97           SFP+-10G-SR
    Xcvr 10      REV 01   740-021308   AQA0GS4           SFP+-10G-SR
    Xcvr 11      REV 01   740-021308   AQA0JVA           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           12X10GE SFPP
    Xcvr 0       REV 01   740-021308   AQA057A           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   ANA0MLS           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQA093A           SFP+-10G-SR
    Xcvr 3       REV 01   740-021309   943153A00075      SFP+-10G-LR
    Xcvr 4       REV 01   740-021308   AQA077B           SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   AQA0JSC           SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AQA0735           SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   AQ5028N           SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   AP40VN5           SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   AQA0K0J           SFP+-10G-SR
    Xcvr 10      REV 01   740-021308   AQA07AP           SFP+-10G-SR
    Xcvr 11      REV 01   740-021308   AQA08YB           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           3X40GE QSFPP
  PIC 3                   BUILTIN      BUILTIN           3X40GE QSFPP
  WAN MEZZ       REV 07   750-045717   CAAX3123          MPC5E 24XGE Mezz
```



```
FPC 11           REV 17   750-037355   CAAT3986         MPC4E 3D 2CGE+8XGE
  CPU            REV 08   711-035209   CAAR3972         HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-021308   AQA0DSE          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQ501Y3          SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQ501XU          SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AQ5036Y          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP
    Xcvr 0                NON-JNPR     X12J00247        CFP-100G-SR10
  PIC 2                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-031980   ALQ1DKF          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJ403YA          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AP40TY0          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   ALQ14G0          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP
    Xcvr 0       REV 01   740-035329   X12J00095        CFP-100G-SR10
Fan Tray 0       REV 08   740-031521   ACAF4219         Enhanced Fan Tray
Fan Tray 1       REV 08   740-031521   ACAF4225         Enhanced Fan Tray
```

show chassis hardware detail (MX960 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                                 MX960
Midplane         REV 01   710-013698   AA6082           MX960 Midplane
PIM              Rev 01   740-013110   000008           Power Inlet Module
PEM 2
PEM 3            Rev 01   740-013682   000038           PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 00   740-015113   1000617944       RE-S-1300
  ad0      245 MB  SanDisk SDCFB-256   111419E1805T1141 Compact Flash
  ad2    38154 MB  FUJITSU MHT2040BH   NR0WT5925N77     Hard Disk
CB 0             REV 05   710-013725   JK6947           MX960 Test SCB
FPC 4            REV 01   710-013305   JM7617           MX960 Test DPC
  CPU
  PIC 0                   BUILTIN      BUILTIN          1x 10GE(LAN/WAN)
  PIC 1                   BUILTIN      BUILTIN          10x 1GE
FPC 7            REV 01   710-013305   JL9634           MX960 Test DPC
  CPU
  PIC 0                   BUILTIN      BUILTIN          1x 10GE(LAN/WAN)
    Xcvr 0                NON-JNPR     MYBG65I82C       XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN          10x 1GE
    Xcvr 1       REV 01   740-011782   P7N0368          SFP-SX
    Xcvr 4       REV 01   740-011782   P8J1W27          SFP-SX
    Xcvr 6       REV 01   740-011782   P8J1VSD          SFP-SX
    Xcvr 9       REV 01   740-011782   P8J1W25          SFP-SX
Fan Tray 0
Fan Tray 1
```

show chassis hardware detail (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN1214852AFA     MX960
Midplane         REV 01   710-030012   ACAX3674         MX960 Backplane
FPM Board        REV 03   710-014974   CAAZ9326         Front Panel Display
PDM             Rev 03   740-013110   QCS17025017     Power Distribution Module
PEM 0            Rev 10   740-027760   QCS1702N062      PS 4.1kW; 200-240V AC
in
```



```
PEM 1           Rev 04   740-027760   QCS1422N02C       PS 4.1kW; 200-240V AC
in
PEM 2           Rev 09   740-027760   QCS1614N01X       PS 4.1kW; 200-240V AC
in
Routing Engine 0 REV 08  740-031116   9009131803        RE-S-1800x4
  ad0    3831 MB  UGB30SFA4000T1      SFA4000T1 000016CD Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI  11000061346       Disk 1
  usb0 (addr 1)  EHCI root hub 0      Intel             uhub0
  usb0 (addr 2)  product 0x0020 32    vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
Routing Engine 1 REV 08  740-031116   9009124913        RE-S-1800x4
  ad0    3831 MB  UGB30SFA4000T1      SFA4000T1 0000106D Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI  11000052402       Disk 1
CB 0            REV 18   750-031391   CABF0579          Enhanced MX SCB
CB 1            REV 16   750-031391   CAAZ2471          Enhanced MX SCB
CB 2            REV 16   750-031391   CAAW9595          Enhanced MX SCB
FPC 0           REV 18   750-046005   CACE6574          MPC5E 3D Q 2CGE+4XGE
  CPU           REV 09   711-045719   CACG8908          RMPC PMB
  PIC 0                  BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0      REV 01   740-021308   AQA0DYT           SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0MS7           SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           1X100GE CFP2 OTN
    Xcvr 0      REV 01   740-046563   XD16FC03Z         CFP2-100G-SR10
  PIC 2                  BUILTIN      BUILTIN           2X10GE SFPP OTN
    Xcvr 0      REV 01   740-021308   ANA0NAJ           SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0MRQ           SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN           1X100GE CFP2 OTN
    Xcvr 0      REV 01   740-049775   J13K72993         CFP2-100G-LR4
FPC 1           REV 11   750-045372   CABK8154          MPCE Type 3 3D
  CPU           REV 08   711-035209   CABE7370          HMPC PMB 2G
  MIC 0         REV 07   750-033307   CABD5255          10X10GE SFPP
    PIC 0                BUILTIN      BUILTIN           10X10GE SFPP
      Xcvr 0    REV 01   740-021308   AQ50319           SFP+-10G-SR
      Xcvr 1    REV 01   740-021308   AQ5035V           SFP+-10G-SR
      Xcvr 2    REV 01   740-021308   AQ502XJ           SFP+-10G-SR
      Xcvr 3    REV 01   740-021308   AQ43HHR           SFP+-10G-SR
      Xcvr 4    REV 01   740-021308   AQ502YA           SFP+-10G-SR
      Xcvr 5    REV 01   740-021308   AQ502EU           SFP+-10G-SR
      Xcvr 6    REV 01   740-021308   AQ502HR           SFP+-10G-SR
      Xcvr 7    REV 01   740-021308   AQ502A6           SFP+-10G-SR
      Xcvr 8    REV 01   740-021308   AQ43H8M           SFP+-10G-SR
  MIC 1         REV 14   750-033196   CAAP1398          1X100GE CXP
    PIC 2                BUILTIN      BUILTIN           1X100GE CXP
      Xcvr 0    REV 01   740-046563   XD16FC064         CFP2-100G-SR10
FPC 3           REV 35   750-028467   CAAT9156          MPC 3D 16x 10GE
  CPU           REV 11   711-029089   CAAV4645          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-021308   AQ43HZ1           SFP+-10G-SR
Xcvr 1      REV 01   740-021308   AQ43HZC           SFP+-10G-SR
    Xcvr 2      REV 01   740-021308   AQ43HD2           SFP+-10G-SR
    Xcvr 3      REV 01   740-021308   AQ502HN           SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-021308   AQ43HGF           SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQ501RZ           SFP+-10G-SR
    Xcvr 2      REV 01   740-021308   AQ5029V           SFP+-10G-SR
    Xcvr 3      REV 01   740-021308   AQ501X9           SFP+-10G-SR
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-021308   AQ502ZN           SFP+-10G-SR
```

```
    Xcvr 1      REV 01   740-021308   AQ43H86          SFP+-10G-SR
    Xcvr 2      REV 01   740-021308   AQ502ZY          SFP+-10G-SR
    Xcvr 3      REV 01   740-021308   AQ502PZ          SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-021308   AQ503E6          SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQ502XN          SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   B11F00213        SFP+-10G-SR
    Xcvr 3      REV 01   740-021308   AQ50336          SFP+-10G-SR
FPC 4           REV 18   750-046005   CACE6568         MPC5E 3D Q 2CGE+4XGE
  CPU           REV 09   711-045719   CACG8900         RMPC PMB
  PIC 0                  BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0      REV 01   740-021308   AQA095A          SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0M1E          SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0               NON-JNPR     FE13F000F        CFP2-100G-SR10
  PIC 2                  BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0      REV 01   740-021308   AQG0LYC          SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0LYB          SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0      REV 01   740-048813   XD32FE00Z        CFP2-100G-SR10
FPC 5           REV 18   750-046005   CACE6577         MPC5E 3D Q 2CGE+4XGE
  CPU           REV 09   711-045719   CACG8902         RMPC PMB
  PIC 0                  BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0      REV 01   740-021308   AQG0MXE          SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0LVY          SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0      REV 01   740-046563   XD16FC03T        CFP2-100G-SR10
  PIC 2                  BUILTIN      BUILTIN          2X10GE SFPP OTN
    Xcvr 0      REV 01   740-021308   AQG0LW1          SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0LW3          SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN          1X100GE CFP2 OTN
    Xcvr 0               NON-JNPR     FE13F000J        CFP2-100G-SR10
FPC 7           REV 09   750-037355   CAAF0937         MPC4E 3D 2CGE+8XGE
  CPU           REV 08   711-035209   CAAD8004         HMPC PMB 2G
  PIC 0                  BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0      REV 01   740-021308   ANA0MM3          SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN          1X100GE CFP
    Xcvr 0      REV 01   740-035329   X000C163         CFP-100G-SR10
  PIC 2                  BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0      REV 01   740-021308   AQG0MS6          SFP+-10G-SR
    Xcvr 1      REV 01   740-021308   AQG0MRX          SFP+-10G-SR
    Xcvr 2      REV 01   740-021308   AQG0M6Y          SFP+-10G-SR
    Xcvr 3      REV 01   740-021308   AQG0LZM          SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN          1X100GE CFP
    Xcvr 0      REV 01   740-035329   X12J00499        CFP-100G-SR10
FPC 8           REV 39   750-045715   CACD1903         MPC5E 3D Q 24XGE+6XLGE
  CPU           REV 09   711-045719   CACD1815         RMPC PMB
  PIC 0                  BUILTIN      BUILTIN          12X10GE SFPP OTN
 PIC 1                  BUILTIN      BUILTIN          12X10GE SFPP OTN
  PIC 2                  BUILTIN      BUILTIN          3X40GE QSFPP
    Xcvr 0      REV 01   740-046565   QC480289         QSFP+-40G-SR4
    Xcvr 1      REV 01   740-046565   QC480274         QSFP+-40G-SR4
    Xcvr 2      REV 01   740-046565   QD130190         QSFP+-40G-SR4
  PIC 3                  BUILTIN      BUILTIN          3X40GE QSFPP
    Xcvr 0      REV 01   740-046565   QD130197         QSFP+-40G-SR4
    Xcvr 1      REV 01   740-046565   QD130180         QSFP+-40G-SR4
    Xcvr 2      REV 01   740-046565   QD130199         QSFP+-40G-SR4
  WAN MEZZ      REV 09   750-049136   CABN0415         MPC5E 24XGE OTN Mezz
FPC 9           REV 05   750-044444   CAAY9801         MPCE Type 2 3D P
  CPU           REV 04   711-038484   CAAW3673         MPCE PMB 2G
  MIC 0         REV 28   750-028387   CAAX1071         3D 4x 10GE  XFP
```



```
    PIC 0                     BUILTIN      BUILTIN           2x 10GE  XFP
      Xcvr 0                  NON-JNPR     T12L92342         XFP-10G-SR
      Xcvr 1                  NON-JNPR     T12L92303         XFP-10G-SR
    PIC 1                     BUILTIN      BUILTIN           2x 10GE  XFP
      Xcvr 0                  NON-JNPR     CC07BK02X         XFP-10G-SR
  QXM 0          REV 06       711-028408   CAAW4883          MPC QXM
  QXM 1          REV 06       711-028408   CAAW4603          MPC QXM
FPC 10           REV 21.0.11 750-045715 CAAY3541           MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 07       711-045719   CAAW7426          RMPC PMB
  PIC 0                       BUILTIN      BUILTIN           12X10GE SFPP
    Xcvr 0       REV 01       740-031980   AHK01AP           SFP+-10G-SR
    Xcvr 1       REV 01       740-021308   AQ502ZU           SFP+-10G-SR
    Xcvr 2       REV 01       740-031980   AP41BLS           SFP+-10G-SR
    Xcvr 3       REV 01       740-021308   AQA08YA           SFP+-10G-SR
    Xcvr 4       REV 01       740-021308   AQA0K26           SFP+-10G-SR
    Xcvr 6       REV 01       740-021308   AQA06S3           SFP+-10G-SR
    Xcvr 7       REV 01       740-021308   AQA06AS           SFP+-10G-SR
    Xcvr 8       REV 01       740-021308   AQA053N           SFP+-10G-SR
    Xcvr 9       REV 01       740-021308   AQA0E97           SFP+-10G-SR
    Xcvr 10      REV 01       740-021308   AQA0GS4           SFP+-10G-SR
    Xcvr 11      REV 01       740-021308   AQA0JVA           SFP+-10G-SR
  PIC 1                       BUILTIN      BUILTIN           12X10GE SFPP
    Xcvr 0       REV 01       740-021308   AQA057A           SFP+-10G-SR
    Xcvr 1       REV 01       740-021308   ANA0MLS           SFP+-10G-SR
    Xcvr 2       REV 01       740-021308   AQA093A           SFP+-10G-SR
    Xcvr 3       REV 01       740-021309   943153A00075      SFP+-10G-LR
    Xcvr 4       REV 01       740-021308   AQA077B           SFP+-10G-SR
    Xcvr 5       REV 01       740-021308   AQA0JSC           SFP+-10G-SR
    Xcvr 6       REV 01       740-021308   AQA0735           SFP+-10G-SR
    Xcvr 7       REV 01       740-021308   AQ5028N           SFP+-10G-SR
    Xcvr 8       REV 01       740-031980   AP40VN5           SFP+-10G-SR
    Xcvr 9       REV 01       740-021308   AQA0K0J           SFP+-10G-SR
    Xcvr 10      REV 01       740-021308   AQA07AP           SFP+-10G-SR
    Xcvr 11      REV 01       740-021308   AQA08YB           SFP+-10G-SR
  PIC 2                       BUILTIN      BUILTIN           3X40GE QSFPP
  PIC 3                       BUILTIN      BUILTIN           3X40GE QSFPP
  WAN MEZZ       REV 07       750-045717   CAAX3123          MPC5E 24XGE Mezz
FPC 11           REV 17       750-037355   CAAT3986          MPC4E 3D 2CGE+8XGE
  CPU            REV 08       711-035209   CAAR3972          HMPC PMB 2G
  PIC 0                       BUILTIN      BUILTIN           4x10GE SFPP
    Xcvr 0       REV 01       740-021308   AQA0DSE           SFP+-10G-SR
    Xcvr 1       REV 01       740-021308   AQ501Y3           SFP+-10G-SR
    Xcvr 2       REV 01       740-021308   AQ501XU           SFP+-10G-SR
 Xcvr 3       REV 01    740-021308   AQ5036Y          SFP+-10G-SR
  PIC 1                       BUILTIN      BUILTIN           1X100GE CFP
    Xcvr 0                    NON-JNPR     X12J00247         CFP-100G-SR10
  PIC 2                       BUILTIN      BUILTIN           4x10GE SFPP
    Xcvr 0       REV 01       740-031980   ALQ1DKF           SFP+-10G-SR
    Xcvr 1       REV 01       740-031980   AJ403YA           SFP+-10G-SR
    Xcvr 2       REV 01       740-031980   AP40TY0           SFP+-10G-SR
    Xcvr 3       REV 01       740-031980   ALQ14G0           SFP+-10G-SR
  PIC 3                       BUILTIN      BUILTIN           1X100GE CFP
    Xcvr 0       REV 01       740-035329   X12J00095         CFP-100G-SR10
Fan Tray 0       REV 08       740-031521   ACAF4219          Enhanced Fan Tray
Fan Tray 1       REV 08       740-031521   ACAF4225          Enhanced Fan Tray
```

show chassis hardware extensive (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
```

```
Chassis                                JN1214852AFA      MX960
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
                                S/N:               JN1214852AFA
 Assembly ID:  0x0512           Assembly Version:  00.00
 Date:         00-00-0000       Assembly Flags:    0x00
 ID: MX960
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 12 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 31 34 38 35 32 41 46 41 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 01   710-030012   ACAX3674          MX960 Backplane
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          710-030012       S/N:               ACAX3674
 Assembly ID:  0x01df           Assembly Version:  01.01
 Date:         01-19-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         COM8T00CRB
 ID: MX960 Backplane            FRU Model Number:  CHAS-BP-MX960-S
 Board Information Record:
  Address 0x00: ad 01 08 00 54 e0 32 bc 68 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 01 df 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 33 30 30 31 32 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 58 33 36 37 34 00 13 01 07
  Address 0x30: dd ff ff ff ad 01 08 00 54 e0 32 bc 68 00 ff ff
  Address 0x40: ff ff ff ff 01 43 4f 4d 38 54 30 30 43 52 42 43
  Address 0x50: 48 41 53 2d 42 50 2d 4d 58 39 36 30 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff aa ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 03   710-014974   CAAZ9326          Front Panel Display
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          710-014974       S/N:               CAAZ9326
 Assembly ID:  0x01e6           Assembly Version:  01.03
 Date:         12-31-2012       Assembly Flags:    0x00
 Version:      REV 03
 ID: Front Panel Display        FRU Model Number:  CRAFT-MX960-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 01 e6 01 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 31 34 39 37 34 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 5a 39 33 32 36 00 1f 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
  Address 0x50: 52 41 46 54 2d 4d 58 39 36 30 2d 53 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PDM              Rev 03   740-013110   QCS17025017       Power Distribution Module
Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          740-013110       S/N:               QCS17025017
 Assembly ID:  0x0416           Assembly Version:  01.03
 Date:         01-10-2013       Assembly Flags:    0x00
 Version:      Rev 03
 ID: Power Distribution Module
 Board Information Record:
```



```
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 16 01 03 52 65 76 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 31 33 31 31 30 00 00
  Address 0x20: 51 43 53 31 37 30 32 35 30 31 37 00 00 0a 01 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 0            Rev 10   740-027760   QCS1702N062       PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-027760        S/N:               QCS1702N062
 Assembly ID:  0x0430            Assembly Version:  01.10
 Date:         01-15-2013        Assembly Flags:    0x00
 Version:      Rev 10
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  PWR-MX960-4100-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 30 01 0a 52 65 76 20 31 30 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
  Address 0x20: 51 43 53 31 37 30 32 4e 30 36 32 00 00 0f 01 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1            Rev 04   740-027760   QCS1422N02C       PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-027760        S/N:               QCS1422N02C
 Assembly ID:  0x0430            Assembly Version:  01.04
 Date:         06-04-2010        Assembly Flags:    0x00
 Version:      Rev 04
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  PWR-MX960-4100-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 30 01 04 52 65 76 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
  Address 0x20: 51 43 53 31 34 32 32 4e 30 32 43 00 00 04 06 07
  Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 2            Rev 09   740-027760   QCS1614N01X       PS 4.1kW; 200-240V AC
in
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          740-027760        S/N:               QCS1614N01X
Assembly ID:  0x0430            Assembly Version:  01.09
 Date:         04-07-2012        Assembly Flags:    0x00
 Version:      Rev 09
 ID: PS 4.1kW; 200-240V AC in    FRU Model Number:  PWR-MX960-4100-AC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 30 01 09 52 65 76 20 30 39 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
```



```
  Address 0x20: 51 43 53 31 36 31 34 4e 30 31 58 00 00 07 04 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 08   740-031116   9009131803        RE-S-1800x4
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-031116       S/N:               9009131803
 Assembly ID:  0x09c0           Assembly Version:  01.08
 Date:         03-04-2013       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         COUCASKBAA
 ID: RE-S-1800x4                FRU Model Number:  RE-S-1800X4-16G-S
 Board Information Record:
  Address 0x00: 54 32 30 32 37 44 42 2d 34 34 47 42 23 42 23 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 31 33 31 38 30 33 00 00 00 04 03 07
  Address 0x30: dd ff ff ff 54 32 30 32 37 44 42 2d 34 34 47 42
  Address 0x40: 23 42 23 00 01 43 4f 55 43 41 53 4b 42 41 41 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 59 ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3831 MB  UGB30SFA4000T1       SFA4000T1 000016CD Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI   11000061346       Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
Routing Engine 1 REV 08   740-031116   9009124913        RE-S-1800x4
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-031116       S/N:               9009124913
 Assembly ID:  0x09c0           Assembly Version:  01.08
 Date:         01-09-2013       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         COUCASKBAA
 ID: RE-S-1800x4                FRU Model Number:  RE-S-1800X4-16G-S
 Board Information Record:
  Address 0x00: 54 32 30 32 37 44 42 2d 34 34 47 42 23 42 23 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 31 32 34 39 31 33 00 00 00 09 01 07
  Address 0x30: dd ff ff ff 54 32 30 32 37 44 42 2d 34 34 47 42
  Address 0x40: 23 42 23 00 01 43 4f 55 43 41 53 4b 42 41 41 52
 Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 59 ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3831 MB  UGB30SFA4000T1       SFA4000T1 0000106D Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI   11000052402       Disk 1
CB 0            REV 18   750-031391   CABF0579          Enhanced MX SCB
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-031391       S/N:               CABF0579
 Assembly ID:  0x09b0           Assembly Version:  01.18
 Date:         04-15-2013       Assembly Flags:    0x00
 Version:      REV 18           CLEI Code:         COUCASRBAA
 ID: Enhanced MX SCB            FRU Model Number:  SCBE-MX-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

```
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 09 b0 01 12 52 45 56 20 31 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
   Address 0x20: 53 2f 4e 20 43 41 42 46 30 35 37 39 00 0f 04 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4f 55 43 41 53 52 42 41 41 53
   Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 43 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 7d ff ff ff ff ff ff ff ff ff ff ff ff
CB 1             REV 16   750-031391   CAAZ2471          Enhanced MX SCB
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-031391        S/N:               CAAZ2471
  Assembly ID:  0x09b0            Assembly Version:  01.16
  Date:         03-09-2013        Assembly Flags:    0x00
  Version:      REV 16            CLEI Code:         COUCARCBAB
  ID: Enhanced MX SCB             FRU Model Number:  SCBE-MX-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 09 b0 01 10 52 45 56 20 31 36 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
   Address 0x20: 53 2f 4e 20 43 41 41 5a 32 34 37 31 00 09 03 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4f 55 43 41 52 43 42 41 42 53
   Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 6d ff ff ff ff ff ff ff ff ff ff ff ff
CB 2             REV 16   750-031391   CAAW9595          Enhanced MX SCB
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-031391        S/N:               CAAW9595
  Assembly ID:  0x09b0            Assembly Version:  01.16
  Date:         02-01-2013        Assembly Flags:    0x00
  Version:      REV 16            CLEI Code:         COUCARCBAB
  ID: Enhanced MX SCB             FRU Model Number:  SCBE-MX-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 09 b0 01 10 52 45 56 20 31 36 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
   Address 0x20: 53 2f 4e 20 43 41 41 57 39 35 39 35 00 01 02 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 43 4f 55 43 41 52 43 42 41 42 53
  Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 6d ff ff ff ff ff ff ff ff ff ff ff ff
FPC 0            REV 18   750-046005   CACE6574          MPC5E 3D Q 2CGE+4XGE
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-046005        S/N:               CACE6574
  Assembly ID:  0x0b8c            Assembly Version:  01.18
  Date:         11-20-2013        Assembly Flags:    0x00
  Version:      REV 18            CLEI Code:         PROTOXCLEI
  ID: MPC5E 3D Q 2CGE+4XGE        FRU Model Number:  PROTO-ASSEMBLY
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 8c 01 12 52 45 56 20 31 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 34 36 30 30 35 00 00
   Address 0x20: 53 2f 4e 20 43 41 43 45 36 35 37 34 00 14 0b 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
   Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
```

```
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
 CPU            REV 09   711-045719   CACG8908          RMPC PMB
Jedec Code:   0x7fb0            EEPROM Version:    0x02
P/N:          711-045719        S/N:               CACG8908
Assembly ID:  0x0b85            Assembly Version:  01.09
Date:         11-13-2013        Assembly Flags:    0x00
Version:      REV 09
ID: RMPC PMB
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 85 01 09 52 45 56 20 30 39 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 34 35 37 31 39 00 00
 Address 0x20: 53 2f 4e 20 43 41 43 47 38 39 30 38 00 0d 0b 07
 Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 00 00 00 00 00 00 00 00 00 00 00 00
 PIC 0                  BUILTIN      BUILTIN           2X10GE SFPP OTN
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0a90            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 2X10GE SFPP OTN
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a 90 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 02 ae dc 00 00 00 00 0a 6e 00 00
Xcvr 0       REV 01   740-021308   AQA0DYT           SFP+-10G-SR
   Xcvr 1       REV 01   740-021308   AQG0MS7           SFP+-10G-SR
 PIC 1                  BUILTIN      BUILTIN           1X100GE CFP2 OTN
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0a6e            Assembly Version:  00.00
Date:         00-00-0000        Assembly Flags:    0x00
ID: 1X100GE CFP2 OTN
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 0a 6e 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 03 f3 8c 31 5c e7 80 00 00 00 02
   Xcvr 0       REV 01   740-046563   XD16FC03Z         CFP2-100G-SR10
 PIC 2                  BUILTIN      BUILTIN           2X10GE SFPP OTN
Jedec Code:   0x0000            EEPROM Version:    0x00
P/N:          BUILTIN           S/N:               BUILTIN
Assembly ID:  0x0a90            Assembly Version:  00.00
```



```
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: 2X10GE SFPP OTN
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 90 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 c0 03 f5 6c 31 5c db 40 00 00 00 02
    Xcvr 0       REV 01   740-021308   ANA0NAJ          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQG0MRQ          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP2 OTN
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0a6e            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: 1X100GE CFP2 OTN
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 6e 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 c0 03 ed ec 31 5c e2 e8 00 00 00 02
 Xcvr 0       REV 01   740-049775   J13K72993         CFP2-100G-LR4
FPC 1            REV 11   750-045372   CABK8154          MPCE Type 3 3D
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-045372        S/N:               CABK8154
 Assembly ID:  0x09db            Assembly Version:  04.11
 Date:         05-18-2013        Assembly Flags:    0x00
 Version:      REV 11            CLEI Code:         COUIBBNBAA
 ID: MPCE Type 3 3D              FRU Model Number:  MX-MPC3E-3D
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 db 04 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 35 33 37 32 00 00
  Address 0x20: 53 2f 4e 20 43 41 42 4b 38 31 35 34 00 12 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4e 42 41 41 4d
  Address 0x50: 58 2d 4d 50 43 33 45 2d 33 44 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 44 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff cf ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 08   711-035209   CABE7370          HMPC PMB 2G
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          711-035209        S/N:               CABE7370
 Assembly ID:  0x0b04            Assembly Version:  01.08
 Date:         05-08-2013        Assembly Flags:    0x00
 Version:      REV 08
 ID: HMPC PMB 2G
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
```

```
  Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
  Address 0x20: 53 2f 4e 20 43 41 42 45 37 33 37 30 00 08 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  MIC 0          REV 07   750-033307   CABD5255          10X10GE SFPP
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-033307       S/N:               CABD5255
 Assembly ID:  0x0a2a           Assembly Version:  02.07
 Date:         04-25-2013       Assembly Flags:    0x00
 Version:      REV 07           CLEI Code:         COUIBBJBAA
 ID: 10X10GE SFPP               FRU Model Number:  MIC3-3D-10XGE-SFPP
 Board Information Record:
  Address 0x00: 34 01 03 03 05 ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0a 2a 02 07 52 45 56 20 30 37 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 33 33 30 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 42 44 35 32 35 35 00 19 04 07
  Address 0x30: dd ff ff ff 34 01 03 03 05 ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4a 42 41 41 4d
  Address 0x50: 49 43 33 2d 33 44 2d 31 30 58 47 45 2d 53 46 50
  Address 0x60: 50 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 82 c0 03 f0 bc 57 79 83 80 00 00 00 02
    PIC 0                 BUILTIN      BUILTIN           10X10GE SFPP
Xcvr 0     REV 01   740-021308   AQ50319           SFP+-10G-SR
      Xcvr 1     REV 01   740-021308   AQ5035V           SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   AQ502XJ           SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   AQ43HHR           SFP+-10G-SR
      Xcvr 4     REV 01   740-021308   AQ502YA           SFP+-10G-SR
      Xcvr 5     REV 01   740-021308   AQ502EU           SFP+-10G-SR
      Xcvr 6     REV 01   740-021308   AQ502HR           SFP+-10G-SR
      Xcvr 7     REV 01   740-021308   AQ502A6           SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AQ43H8M           SFP+-10G-SR
  MIC 1          REV 14   750-033196   CAAP1398          1X100GE CXP
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-033196       S/N:               CAAP1398
 Assembly ID:  0x0a29           Assembly Version:  03.14
 Date:         10-27-2012       Assembly Flags:    0x00
 Version:      REV 14           CLEI Code:         COUIBBKBAA
 ID: 1X100GE CXP                FRU Model Number:  MIC3-3D-1X100GE-CXP
 Board Information Record:
  Address 0x00: 34 01 07 07 08 ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0a 29 03 0e 52 45 56 20 31 34 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 33 31 39 36 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 50 31 33 39 38 00 1b 0a 07
  Address 0x30: dc ff ff ff 34 01 07 07 08 ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4b 42 41 41 4d
  Address 0x50: 49 43 33 2d 33 44 2d 31 58 31 30 30 47 45 2d 43
  Address 0x60: 58 50 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 96 c0 03 ef cc 57 79 85 08 00 00 00 02
    PIC 2                 BUILTIN      BUILTIN           1X100GE CXP
      Xcvr 0     REV 01   740-046563   XD16FC064         CFP2-100G-SR10
FPC 3            REV 35   750-028467   CAAT9156          MPC 3D 16x 10GE
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          750-028467       S/N:               CAAT9156
 Assembly ID:  0x0997           Assembly Version:  01.35
 Date:         12-17-2012       Assembly Flags:    0x00
```

```
Version:      REV 35
ID: MPC 3D 16x 10GE            FRU Model Number:  MPC-3D-16XGE-SFPP
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 09 97 01 23 52 45 56 20 33 35 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 32 38 34 36 37 00 00
 Address 0x20: 53 2f 4e 20 43 41 41 54 39 31 35 36 00 11 0c 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 4d
 Address 0x50: 50 43 2d 33 44 2d 31 36 58 47 45 2d 53 46 50 50
 Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 CPU           REV 11   711-029089   CAAV4645          AMPC PMB
Jedec Code:   0x7fb0           EEPROM Version:    0x01
P/N:          711-029089       S/N:               CAAV4645
Assembly ID:  0x0998           Assembly Version:  01.11
Date:         12-13-2012       Assembly Flags:    0x00
Version:      REV 11
ID: AMPC PMB
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 09 98 01 0b 52 45 56 20 31 31 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 32 39 30 38 39 00 00
 Address 0x20: 53 2f 4e 20 43 41 41 56 34 36 34 35 00 0d 0c 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
 PIC 0                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
Jedec Code:   0x0000           EEPROM Version:    0x00
P/N:          BUILTIN          S/N:               BUILTIN
Assembly ID:  0x02fe           Assembly Version:  00.00
Date:         00-00-0000       Assembly Flags:    0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
 Address 0x00: 00 00 00 00 02 fe 00 00 00 00 00 00 00 00 00 00
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 Address 0x70: 00 00 00 00 c0 02 6b 94 00 00 00 00 02 fe 00 00
   Xcvr 0       REV 01   740-021308   AQ43HZ1          SFP+-10G-SR
   Xcvr 1       REV 01   740-021308   AQ43HZC          SFP+-10G-SR
   Xcvr 2       REV 01   740-021308   AQ43HD2          SFP+-10G-SR
   Xcvr 3       REV 01   740-021308   AQ502HN          SFP+-10G-SR
 PIC 1                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
Jedec Code:   0x0000           EEPROM Version:    0x00
P/N:          BUILTIN          S/N:               BUILTIN
Assembly ID:  0x02fe           Assembly Version:  00.00
Date:         00-00-0000       Assembly Flags:    0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
 Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
```

```
  Address 0x00: 00 00 00 00 02 fe 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
  Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 c0 02 ac 0c 00 00 00 00 02 fe 00 00
    Xcvr 0        REV 01   740-021308   AQ43HGF          SFP+-10G-SR
    Xcvr 1        REV 01   740-021308   AQ501RZ          SFP+-10G-SR
    Xcvr 2        REV 01   740-021308   AQ5029V          SFP+-10G-SR
    Xcvr 3        REV 01   740-021308   AQ501X9          SFP+-10G-SR
  PIC 2                    BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x02fe            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
......
```

show chassis hardware models (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware models
Hardware inventory:
Item                Version   Part number  Serial number     FRU model number
Midplane            REV 01    710-030012   ACAX3674          CHAS-BP-MX960-S
FPM Board           REV 03    710-014974   CAAZ9326          CRAFT-MX960-S
PEM 0               Rev 10    740-027760   QCS1702N062       PWR-MX960-4100-AC-S
PEM 1               Rev 04    740-027760   QCS1422N02C       PWR-MX960-4100-AC-S
PEM 2               Rev 09    740-027760   QCS1614N01X       PWR-MX960-4100-AC-S
Routing Engine 0 REV 08       740-031116   9009131803        RE-S-1800X4-16G-S
Routing Engine 1 REV 08       740-031116   9009124913        RE-S-1800X4-16G-S
CB 0                REV 18    750-031391   CABF0579          SCBE-MX-S
CB 1                REV 16    750-031391   CAAZ2471          SCBE-MX-S
CB 2                REV 16    750-031391   CAAW9595          SCBE-MX-S
FPC 0               REV 18    750-046005   CACE6574          PROTO-ASSEMBLY
FPC 1               REV 11    750-045372   CABK8154          MX-MPC3E-3D
  MIC 0             REV 07    750-033307   CABD5255          MIC3-3D-10XGE-SFPP
  MIC 1             REV 14    750-033196   CAAP1398          MIC3-3D-1X100GE-CXP
FPC 3               REV 35    750-028467   CAAT9156          MPC-3D-16XGE-SFPP
FPC 4               REV 18    750-046005   CACE6568          PROTO-ASSEMBLY
FPC 5               REV 18    750-046005   CACE6577          PROTO-ASSEMBLY
FPC 7               REV 09    750-037355   CAAF0937          MPC4E-2CGE-8XGE
FPC 8               REV 39    750-045715   CACD1903          PROTO-ASSEMBLY
FPC 9               REV 05    750-044444   CAAY9801          MX-MPC2E-3D-P
  MIC 0             REV 28    750-028387   CAAX1071          MIC-3D-4XGE-XFP
FPC 10              REV 21.0.11 750-045715 CAAY3541          PROTO-ASSEMBLY
FPC 11              REV 17    750-037355   CAAT3986          MPC4E-3D-2CGE-8XGE
Fan Tray 0          REV 08    740-031521   ACAF4219          FFANTRAY-MX960-HC-S
Fan Tray 1          REV 08    740-031521   ACAF4225          FFANTRAY-MX960-HC-S
```

show chassis hardware clei-models (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version   Part number  CLEI code         FRU model number
Midplane            REV 01    710-030012   COM8T00CRB        CHAS-BP-MX960-S
FPM Board           REV 03    710-014974                     CRAFT-MX960-S
PEM 0               Rev 10    740-027760                     PWR-MX960-4100-AC-S
PEM 1               Rev 04    740-027760                     PWR-MX960-4100-AC-S
PEM 2               Rev 09    740-027760                     PWR-MX960-4100-AC-S
```

```
Routing Engine 0 REV 08    740-031116    COUCASKBAA         RE-S-1800X4-16G-S
Routing Engine 1 REV 08    740-031116    COUCASKBAA         RE-S-1800X4-16G-S
CB 0             REV 18    750-031391    COUCASRBAA         SCBE-MX-S
CB 1             REV 16    750-031391    COUCARCBAB         SCBE-MX-S
CB 2             REV 16    750-031391    COUCARCBAB         SCBE-MX-S
FPC 0            REV 18    750-046005    PROTOXCLEI         PROTO-ASSEMBLY
FPC 1            REV 11    750-045372    COUIBBNBAA         MX-MPC3E-3D
  MIC 0          REV 07    750-033307    COUIBBJBAA         MIC3-3D-10XGE-SFPP
  MIC 1          REV 14    750-033196    COUIBBKBAA         MIC3-3D-1X100GE-CXP
FPC 3            REV 35    750-028467                       MPC-3D-16XGE-SFPP
FPC 4            REV 18    750-046005    PROTOXCLEI         PROTO-ASSEMBLY
FPC 5            REV 18    750-046005    PROTOXCLEI         PROTO-ASSEMBLY
FPC 7            REV 09    750-037355    PROTOXCLEI         MPC4E-2CGE-8XGE
FPC 8            REV 39    750-045715    PROTOXCLEI         PROTO-ASSEMBLY
FPC 9            REV 05    750-044444    COUIBBGBAA         MX-MPC2E-3D-P
  MIC 0          REV 28    750-028387    COUIA16BAA         MIC-3D-4XGE-XFP
FPC 10           REV 21.0.11 750-045715  PROTOXCLEI         PROTO-ASSEMBLY
FPC 11           REV 17    750-037355    IPU3A4DHAA         MPC4E-3D-2CGE-8XGE
Fan Tray 0       REV 08    740-031521                       FFANTRAY-MX960-HC-S
Fan Tray 1       REV 08    740-031521                       FFANTRAY-MX960-HC-S
```

show chassis hardware (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN123F6D9AFA      MX960
Midplane         REV 04   750-047849   ACRC8764          Enhanced MX960 Backplane
FPM Board        REV 03   710-014974   CACS4395          Front Panel Display
PDM              Rev 03   740-013110   QCS1809500Z       Power Distribution Module
PEM 0            Rev 08   740-029344   QCS1817V0LK       DC 4.1kW Power Entry
Module
PEM 1            Rev 08   740-029344   QCS1814V01F       DC 4.1kW Power Entry
Module
PEM 2            Rev 08   740-029344   QCS1810V1EW       DC 4.1kW Power Entry
Module
PEM 3            Rev 08   740-029344   QCS1810V1K5       DC 4.1kW Power Entry
Module
Routing Engine 0 REV 11   740-031116   9013103483        RE-S-1800x4
Routing Engine 1 REV 10   740-031116   9009198513        RE-S-1800x4
CB 0             REV 23   750-031391   CADW3218          Enhanced MX SCB
CB 1             REV 14   750-031391   ABBK5220          Enhanced MX SCB
FPC 1            REV 14   750-045372   CADK0464          MPCE Type 3 3D
  CPU            REV 10   711-035209   CADM9839          HMPC PMB 2G
  MIC 0          REV 19   750-033199   CAAE5870          1X100GE CFP
    PIC 0                 BUILTIN      BUILTIN           1X100GE CFP
      Xcvr 0     REV 01   740-032210   UTH0H0W           CFP-100G-LR4
FPC 2            REV 14   750-045372   CADN3262          MPCE Type 3 3D
  CPU            REV 10   711-035209   CADN8129          HMPC PMB 2G
FPC 3            REV 14   750-045372   CADH0146          MPCE Type 3 3D
  CPU            REV 10   711-035209   CADT2458          HMPC PMB 2G
  MIC 0          REV 03   750-057666   CADP1386          1X100GE DWDM CFP2-ACO
    PIC 0                 BUILTIN      BUILTIN           1X100GE DWDM CFP2-ACO
      Xcvr 0     REV 01   740-062357   SMD5136.1         OTN-100G-LH
FPC 4            REV 18   750-045372   CAEV5668          MPCE Type 3 3D
  CPU            REV 10   711-035209   CAET7827          HMPC PMB 2G
FPC 7            REV 14   750-045372   CADJ1947          MPCE Type 3 3D
  CPU            REV 10   711-035209   CADJ1561          HMPC PMB 2G
  MIC 0          REV 05   750-057666   CAEB5763          1X100GE DWDM CFP2-ACO
    PIC 0                 BUILTIN      BUILTIN           1X100GE DWDM CFP2-ACO
      Xcvr 0     REV 01   740-062357   1DJBZ052002       OTN-100G-LH
```

```
FPC 8            REV 14   750-045372   CADK0485          MPCE Type 3 3D
  CPU            REV 10   711-035209   CADM9828          HMPC PMB 2G
  MIC 0          REV 03   750-057666   CADP1390          1X100GE DWDM CFP2-ACO
    PIC 0                 BUILTIN      BUILTIN           1X100GE DWDM CFP2-ACO
FPC 9            REV 14   750-045372   CADJ1936          MPCE Type 3 3D
  CPU            REV 10   711-035209   CADJ1566          HMPC PMB 2G
  MIC 0          REV 14   750-057666   CAFF7544          1X100GE DWDM CFP2-ACO
    PIC 0                 BUILTIN      BUILTIN           1X100GE DWDM CFP2-ACO
      Xcvr 0     REV 01   740-062357   1DJBZ05100K       OTN-100G-LH
FPC 10           REV 14   750-054901   CADJ3846          MPC3E NG HQoS
  CPU            REV 11   711-045719   CADN5471          RMPC PMB
  MIC 0          REV 05   750-057666   CAEB5760          1X100GE DWDM CFP2-ACO
 PIC 0                 BUILTIN      BUILTIN           1X100GE DWDM CFP2-ACO
      Xcvr 0     REV 01   740-062357   SMD5091.1         CFP-Loopback
Fan Tray 0       REV 08   740-031521   ACDB4083          Enhanced Fan Tray
Fan Tray 1       REV 08   740-031521   ACDB3995          Enhanced Fan Tray
```

show chassis hardware clei-models(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 04   750-047849   CMMJA10BRA        CHAS-BP3-MX960-S
FPM Board        REV 03   710-014974                     CRAFT-MX960-S
PEM 0            Rev 08   740-029344                     PWR-MX960-4100-DC-S
PEM 1            Rev 08   740-029344                     PWR-MX960-4100-DC-S
PEM 2            Rev 08   740-029344                     PWR-MX960-4100-DC-S
PEM 3            Rev 08   740-029344                     PWR-MX960-4100-DC-S
Routing Engine 0 REV 11   740-031116   COUCASYBAB        RE-S-1800X4-16G-S
Routing Engine 1 REV 10   740-031116   COUCASYBAA        RE-S-1800X4-16G-S
CB 0             REV 23   750-031391   COUCATXBAA        SCBE-MX-S
CB 1             REV 14   750-031391   COUCARCBAA        SCBE-MX-S
FPC 1            REV 14   750-045372   COUIBBNBAB        MX-MPC3E-3D
  MIC 0          REV 19   750-033199   COUIBA8BAA        MIC3-3D-1X100GE-CFP
FPC 2            REV 14   750-045372   COUIBBNBAB        MX-MPC3E-3D
FPC 3            REV 14   750-045372   COUIBBNBAB        MX-MPC3E-3D
  MIC 0          REV 03   750-057666   PROTOXCLEI        PROTO-ASSEMBLY
FPC 4            REV 18   750-045372   COUIBBNBAC        MX-MPC3E-3D
FPC 7            REV 14   750-045372   COUIBBNBAB        MX-MPC3E-3D
  MIC 0          REV 05   750-057666   PROTOXCLEI        PROTO-ASSEMBLY
FPC 8            REV 14   750-045372   COUIBBNBAB        MX-MPC3E-3D
  MIC 0          REV 03   750-057666   PROTOXCLEI        PROTO-ASSEMBLY
FPC 9            REV 14   750-045372   COUIBBNBAB        MX-MPC3E-3D
  MIC 0          REV 14   750-057666   PROTOXCLEI        PROTO-ASSEMBLY
FPC 10           REV 14   750-054901   PROTOXCLEI        PROTO-ASSEMBLY
  MIC 0          REV 05   750-057666   PROTOXCLEI        PROTO-ASSEMBLY
Fan Tray 0       REV 08   740-031521                     FFANTRAY-MX960-HC-S
Fan Tray 1       REV 08   740-031521                     FFANTRAY-MX960-HC-S
```

show chassis hardware (MX10008 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BLANK             JNP10008-MX
Midplane         REV 03   750-054097   ACAM1747          JNP10008 Midplane
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x10
Routing Engine 1          BUILTIN      BUILTIN           RE-S-2X00x10
CB 0             REV 17   750-052688   ACAM7468          Control Board
CB 1             REV 18   750-052688   ACAM7709          Control Board
```

```
FPC 0            REV 33   750-054576   CAFC8443         MPCAE 3D
  CPU                     BUILTIN      BUILTIN          MPC CPU
  PIC 0                   BUILTIN      BUILTIN          MRATE-8xQSFPP-XGE-XLGE
    Xcvr 0       REV 01   740-032986   QD472831         QSFP+-40G-SR4
    Xcvr 1       REV 01   740-032986   QD472839         QSFP+-40G-SR4
    Xcvr 2       REV 01   740-032986   QB421310         QSFP+-40G-SR4
    Xcvr 3       REV 01   740-032986   QD472831         QSFP+-40G-SR4
  .
.
.


    PIC 5                 BUILTIN      BUILTIN
MRATE-16xQSFPP-XGE-XLGE-CGE
      Xcvr 0     REV 01   740-054053   QE419464         QSFP+-4X10G-SR
      Xcvr 1     REV 01   740-046565   QE413929         QSFP+-40G-SR4
      Xcvr 2     REV 01   740-058734   1ACQ1042028      QSFP28-100GBASE-SR4
      Xcvr 3     REV 01   740-046565   QE414116         QSFP+-40G-SR4
FPC 1            REV 33   750-054576   CAFC8443         MPCCE 3D
  CPU                     BUILTIN      BUILTIN          MPC CPU
  PIC 0                   BUILTIN      BUILTIN          MRATE-8xQSFPP-XGE-XLGE
    Xcvr 0       REV 01   740-032986   QD472831         QSFP+-40G-SR4
    Xcvr 2       REV 01   740-032986   QB421310         QSFP+-40G-SR4
    Xcvr 4       REV 01   740-032986   QA480242         QSFP+-40G-SR4
    Xcvr 5       REV 01   740-032986   QE112585         QSFP+-40G-SR4
FPD Board        REV 01   740-XXXXXX   XXXXXXXX         Front Panel Display
Power Supply 0   REV 02   740-049388   1EDL534003N      AC 2850W Power Supply
Power Supply 1   REV 01   740-049388   1EDL44300CF      AC 2850W Power Supply
Power Supply 2   REV 02   740-049388   1EDL534004F      AC 2850W Power Supply
Power Supply 3   REV 02   740-049388   1EDL5340049      AC 2850W Power Supply
FTC 0            REV 08   750-050108   ACAM7310         Fan Tray Controller
FTC 1            REV 08   750-050108   ACAM7316         Fan Tray Controller
Fan Tray 0       REV 01                ACAM1683         Top Fan Tray
Fan Tray 1       REV 01   760-054372   ACAM1657         Vertical Fan Tray
SFB 0            REV 13   750-050058   ACAM8990         Switch Fabric Board
SFB 1            REV 13   750-050058   ACAM8978         Switch Fabric Board
SFB 2            REV 10   750-050058   ACAM8350         Switch Fabric Board
SFB 3            REV 10   750-050058   ACAM8365         Switch Fabric Board
SFB 4            REV 13   750-050058   ACAM8941         Switch Fabric Board
SFB 5            REV 13   750-050058   ACAM8925         Switch Fabric Board
```

show chassis hardware (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN123AC42AJC     PTX3000
Midplane         REV 22   750-044645   ACLP6640         Backplane
FPM              REV 07   760-044663   ACMX2146         Front Panel Display
PSM 1            REV 02   740-044980   1EDD3080169      DC 12V Power Supply
PSM 2            REV 06   740-044981   1EDK5040563      AC 12V Power Supply
PSM 3            REV 06   740-044981   1EDK5040313      AC 12V Power Supply
PSM 4            REV 04   740-044980   1EDJ3330088      DC 12V Power Supply
Routing Engine 0 REV 12   740-026942   P737A-006029     RE-DUO-2600
CB 0             REV 18   750-044656   ACMZ3179         Control Board
FPC 2            REV 06   750-057064   ACAM6098         FPC3-SFF-PTX-1X
  CPU                     BUILTIN      BUILTIN          SMPC PMB
  PIC 0          REV 17   750-059747   ACNW3510         5X100GE DWDM CFP2-ACO
    Xcvr 0       REV 01   740-062357   1DJBZ040003      OTN-100G-LH
    Xcvr 2       REV 01   740-062357   1DJBZ044004      OTN-100G-LH
    Xcvr 3       REV 01   740-062357   1DJBZ03500P      OTN-100G-LH
```

```
    Xcvr 4          REV 01   740-062357   1DJBZ03700C       OTN-100G-LH
FPC 4               REV 12   750-057064   ACAM7153          FPC3-SFF-PTX-1X
  CPU                        BUILTIN      BUILTIN           SMPC PMB
  PIC 0             REV 17   750-059747   ACNW3511          5X100GE DWDM CFP2-ACO
    Xcvr 0          REV 01   740-061663   47                OTN-100G-LH
    Xcvr 1          REV 01   740-061663   39                OTN-100G-LH
    Xcvr 2          REV 01   740-062357   1DJBZ044002       OTN-100G-LH
    Xcvr 3          REV 01   740-062357   1DJBZ03700G       OTN-100G-LH
    Xcvr 4          REV 01   740-062357   1DJBZ041001       OTN-100G-LH
FPC 8               REV 11   750-057064   ACAM6808          FPC3-SFF-PTX-1X
  CPU                        BUILTIN      BUILTIN           SMPC PMB
  PIC 0             REV 17   750-059747   ACNW3508          5X100GE DWDM CFP2-ACO
    Xcvr 0          REV 01   740-061663   194               OTN-100G-LH
    Xcvr 1          REV 01   740-061663   168               OTN-100G-LH
    Xcvr 2          REV 01   740-061663   52                OTN-100G-LH
    Xcvr 3          REV 01   740-061663   85                OTN-100G-LH
    Xcvr 4          REV 01   740-061663   218               OTN-100G-LH
SIB 0               REV 03   750-057067   ACAM8513          SIB3-SFF-PTX
SIB 1               REV 01   750-057067   ACAM5918          SIB3-SFF-PTX
SIB 2               REV 01   711-057066   ACAM4325          SIB3-SFF-PTX
SIB 3               REV 01   711-057066   ACAM4328          SIB3-SFF-PTX
SIB 4               REV 01   711-057066   ACAM4349          SIB3-SFF-PTX
SIB 5               REV 01   711-057066   ACAM4323          SIB3-SFF-PTX
SIB 6               REV 01   711-057066   ACAM4344          SIB3-SFF-PTX
SIB 7               REV 01   750-057067   ACAM4346          SIB3-SFF-PTX
SIB 8               REV 01   750-057067   ACAM5911          SIB3-SFF-PTX
Fan Tray 0          REV 13   760-044659   ACMP6395          Fan Tray (Exhaust)
Fan Tray 1          REV 13   760-044659   ACMZ6957          Fan Tray (Exhaust)
```

show chassis hardware clei-models (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code         FRU model number
Midplane            REV 22   750-044645   IPMVN10FRA        CHAS-MP-PTX3000-S
FPM                 REV 07   760-044663   IPUCBE5CAA        FPD-SFF-PTX-S
PSM 1               REV 02   740-044980   PROTOPWRDC        PSM-SFF-PTX-DC-2200-S
PSM 2               REV 06   740-044981   IPUPAK0KAB        PSM-SFF-PTX-AC-S
PSM 3               REV 06   740-044981   IPUPAK0KAB        PSM-SFF-PTX-AC-S
PSM 4               REV 04   740-044980   IPUPAK1KAA        PSM-SFF-PTX-DC-S
Routing Engine 0    REV 12   740-026942                     RE-DUO-C2600-16G-S
CB 0                REV 18   750-044656   IPUCBE6CAB        CB-SFF-PTX-S
FPC 2               REV 06   750-057064   PROTOXCLEI        PROTO-ASSEMBLY
  PIC 0             REV 17   750-059747   IPU3BC5HAA        PTX-5-100G-WDM
FPC 4               REV 12   750-057064
  PIC 0             REV 17   750-059747   IPU3BC5HAA        PTX-5-100G-WDM
FPC 8               REV 11   750-057064
  PIC 0             REV 17   750-059747   IPU3BC5HAA        PTX-5-100G-WDM
SIB 0               REV 03   750-057067   PROTOXCLEI        PROTO-ASSEMBLY
SIB 1               REV 01   750-057067   PROTOXCLEI        PROTO-ASSEMBLY
SIB 2               REV 01   711-057066   PROTOXCLEI        PROTO-ASSEMBLY
SIB 3               REV 01   711-057066   PROTOXCLEI        PROTO-ASSEMBLY
SIB 4               REV 01   711-057066   PROTOXCLEI        PROTO-ASSEMBLY
SIB 5               REV 01   711-057066   PROTOXCLEI        PROTO-ASSEMBLY
SIB 6               REV 01   711-057066   PROTOXCLEI        PROTO-ASSEMBLY
SIB 7               REV 01   750-057067   PROTOXCLEI        PROTO-ASSEMBLY
SIB 8               REV 01   750-057067   PROTOXCLEI        PROTO-ASSEMBLY
Fan Tray 0          REV 13   760-044659   IPUCBE8CAA        FAN-SFF-PTX-S
Fan Tray 1          REV 13   760-044659   IPUCBE8CAA        FAN-SFF-PTX-S
```


show chassis hardware (MX2010 Router)

```
user@host > show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11E3217AFK      MX2010
Midplane         REV 01   750-044636   ABAB8506          Lower Backplane
Midplane 1       REV 01   711-044557   ZY8296            Upper Backplane
PMP              REV 03   711-032426   ACAJ1388          Power Midplane
FPM Board        REV 06   711-032349   ZX8744            Front Panel Display
PSM 4            REV 0C   740-033727   VK00254           DC 52V Power Supply
Module
PSM 5            REV 0B   740-033727   VG00015           DC 52V Power Supply
Module
PSM 6            REV 0B   740-033727   VH00097           DC 52V Power Supply
Module
PSM 7            REV 0C   740-033727   VJ00151           DC 52V Power Supply
Module
PSM 8            REV 0C   740-033727   VJ00149           DC 52V Power Supply
Module
PDM 0            REV 0B   740-038109   WA00008           DC Power Dist Module
PDM 1            REV 0B   740-038109   WA00014           DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009094134        RE-S-1800x4
Routing Engine 1 REV 02   740-041821   9009094141        RE-S-1800x4
CB 0             REV 08   750-040257   CAAB3491          Control Board
CB 1             REV 08   750-040257   CAAB3489          Control Board
SPMB 0           REV 02   711-041855   CAAA6135          PMB Board
SPMB 1           REV 02   711-041855   CAAA6137          PMB Board
SFB 0            REV 06   711-032385   ZV1828            Switch Fabric Board
SFB 1            REV 07   711-032385   ZZ2568            Switch Fabric Board
SFB 2            REV 07   711-032385   ZZ2563            Switch Fabric Board
SFB 3            REV 07   711-032385   ZZ2564            Switch Fabric Board
SFB 4            REV 07   711-032385   ZZ2580            Switch Fabric Board
SFB 5            REV 07   711-032385   ZZ2579            Switch Fabric Board
SFB 6            REV 07   711-032385   CAAB4882          Switch Fabric Board
SFB 7            REV 07   711-032385   CAAB4898          Switch Fabric Board
FPC 0            REV 33   750-028467   CAAB1919          MPC 3D 16x 10GE
  CPU            REV 11   711-029089   CAAB7174          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AMH02RE           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AMH038C           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AMH0390           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AMG0SUA           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AMH0579           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AMG0SGP           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AMH04SV           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AMH04X3           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AMH0135           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AMH02NC           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AMH02XB           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AMH02PN           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AMH057Y           SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AMG0JHE           SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AMH02HT           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AMH04V4           SFP+-10G-SR
FPC 1            REV 21   750-033205   ZG5027            MPC Type 3
  CPU            REV 04   711-035209   YT4780            HMPC PMB 2G
  MIC 0          REV 03   750-033307   ZV6299            10X10GE SFPP
```



```
    PIC 0                 BUILTIN      BUILTIN           10X10GE SFPP
      Xcvr 0     REV 01   740-031980   083363A00410      SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   083363A00334      SFP+-10G-SR
      Xcvr 2     REV 01   740-031980   113363A00125      SFP+-10G-SR
      Xcvr 3     REV 01   740-031980   083363A00953      SFP+-10G-SR
      Xcvr 4     REV 01   740-031980   AHR013D           SFP+-10G-SR
      Xcvr 5     REV 01   740-031980   AJ40JUR           SFP+-10G-SR
      Xcvr 6     REV 01   740-031980   AJ40JKL           SFP+-10G-SR
      Xcvr 7     REV 01   740-031980   AJ30ECK           SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   19T511100864      SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   19T511100868      SFP+-10G-SR
  MIC 1          REV 03   750-033307   ZV6268            10X10GE SFPP
    PIC 2                 BUILTIN      BUILTIN           10X10GE SFPP
      Xcvr 0     REV 01   740-031980   AJC0JML           SFP+-10G-SR
      Xcvr 1     REV 01   740-031980   AJ403PC           SFP+-10G-SR
      Xcvr 2     REV 01   740-031980   AJ10N25           SFP+-10G-SR
      Xcvr 3     REV 01   740-031980   AJ40JF4           SFP+-10G-SR
      Xcvr 4     REV 01   740-031980   AJ40JSJ           SFP+-10G-SR
      Xcvr 5     REV 01   740-031980   AJ403V7           SFP+-10G-SR
      Xcvr 6     REV 01   740-031980   AJ40JN3           SFP+-10G-SR
      Xcvr 7     REV 01   740-031980   AJ40JSU           SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   19T511100468      SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   19T511101363      SFP+-10G-SR
FPC 8            REV 22   750-031089   ZT9746            MPC Type 2 3D
  CPU            REV 06   711-030884   ZS1271            MPC PMB 2G
  MIC 0          REV 26   750-028392   ABBS1150          3D 20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   PLG023C           SFP-SX
      Xcvr 1     REV 01   740-031851   PLG09C6           SFP-SX
      Xcvr 2     REV 02   740-011613   AM0950SF9L7       SFP-SX
      Xcvr 3     REV 02   740-011613   AM1001SFN1H       SFP-SX
      Xcvr 4     REV 02   740-011613   AM1001SFM9D       SFP-SX
      Xcvr 5     REV 02   740-011613   AM1001SFLTJ       SFP-SX
      Xcvr 6     REV 01   740-031851   AC1108S03L9       SFP-SX
      Xcvr 7     REV 01   740-031851   AC1102S00NC       SFP-SX
      Xcvr 8     REV 01   740-031851   AC1102S00MX       SFP-SX
      Xcvr 9     REV 01   740-031851   AC1102S0085       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   AC1102S00KU       SFP-SX
      Xcvr 1     REV 01   740-031851   AC1102S00NG       SFP-SX
      Xcvr 2     REV 01   740-031851   AC1102S00K3       SFP-SX
      Xcvr 3     REV 01   740-031851   AC1102S008R       SFP-SX
      Xcvr 4     REV 01   740-031851   AM1107SUFVJ       SFP-SX
      Xcvr 5     REV 01   740-031851   AC1108S03LG       SFP-SX
  MIC 1          REV 26   750-028387   ABBR9582          3D 4x 10GE  XFP
    PIC 2                 BUILTIN      BUILTIN           2x 10GE  XFP
      Xcvr 0              NON-JNPR     T10A91703         XFP-10G-SR
      Xcvr 1              NON-JNPR     T09L42604         XFP-10G-SR
    PIC 3                 BUILTIN      BUILTIN           2x 10GE  XFP
FPC 9            REV 11   750-036284   ZL3591            MPC 3D 16x 10GE EM
  CPU            REV 10   711-029089   ZL0513            AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   1YT517101825      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   1YT517101821      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   1YT517101682      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   ALQ13R6           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   1YT517101828      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   1YT517101716      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   1YT517101732      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   ALP0TR1           SFP+-10G-SR
```



```
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   1YT517101741      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   1YT517101829      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   1YT517101669      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   ALQ14E3           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   1YT517101826      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   1YT517101817      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   1YT517101735      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   ALQ159A           SFP+-10G-SR
ADC 0            REV 05   750-043596   CAAC2073          Adapter Card
ADC 1            REV 01   750-043596   ZV4117            Adapter Card
ADC 8            REV 01   750-043596   ZV4107            Adapter Card
ADC 9            REV 02   750-043596   ZW1555            Adapter Card
Fan Tray 0       REV 2A   760-046960   ACAY0015          172mm FanTray - 6 Fans
Fan Tray 1       REV 2A   760-046960   ACAY0019          172mm FanTray - 6 Fans
Fan Tray 2       REV 2A   760-046960   ACAY0020          172mm FanTray - 6 Fans
Fan Tray 3       REV 2A   760-046960   ACAY0021          172mm FanTray - 6 Fans
```

show chassis hardware detail (MX2010 Router)

```
user@host > show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11E233DAFK      MX2010
Midplane         REV 26   750-044636   ABAB9357          Lower Backplane
Midplane 1       REV 01   711-044557   ABAB8643          Upper Backplane
PMP              REV 04   711-032426   ACAJ1677          Power Midplane
FPM Board        REV 08   760-044634   ABBV9726          Front Panel Display
PSM 0            REV 01   740-045050   1E02224000P       DC 52V Power Supply
Module
PSM 1            REV 01   740-045050   1E02224000M       DC 52V Power Supply
Module
PSM 2            REV 01   740-045050   1E022240010       DC 52V Power Supply
Module
PSM 3            REV 01   740-045050   1E02224000G       DC 52V Power Supply
Module
PSM 4            REV 01   740-045050   1E022240013       DC 52V Power Supply
Module
PSM 5            REV 01   740-045050   1E022240007       DC 52V Power Supply
Module
PSM 6            REV 01   740-045050   1E02224001C       DC 52V Power Supply
Module
PSM 7            REV 01   740-045050   1E02224001D       DC 52V Power Supply
Module
PSM 8            REV 01   740-045050   1E02224001B       DC 52V Power Supply
Module
PDM 0            REV 01   740-045234   1E262250067       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009099704        RE-S-1800x4
  ad0    3831 MB  UGB30SFA4000T1       SFA4000T1 00000651 Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI   11000019592       Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009099706        RE-S-1800x4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-499551-000404 Disk 1
CB 0             REV 13   750-040257   CAAF8436          Control Board
```

```
CB 1             REV 13   750-040257   CAAF8434         Control Board
SPMB 0           REV 02   711-041855   ABBV3825         PMB Board
SPMB 1           REV 02   711-041855   ABBV3833         PMB Board
SFB 0            REV 05   711-044466   ABBX5682         Switch Fabric Board
SFB 1            REV 05   711-044466   ABBX5676         Switch Fabric Board
SFB 2            REV 05   711-044466   ABBX5665         Switch Fabric Board
SFB 3            REV 05   711-044466   ABBX5699         Switch Fabric Board
SFB 4            REV 05   711-044466   ABBX5603         Switch Fabric Board
SFB 5            REV 05   711-044466   ABBX5587         Switch Fabric Board
SFB 6            REV 05   711-044466   ABBX5607         Switch Fabric Board
SFB 7            REV 05   711-044466   ABBX5669         Switch Fabric Board
FPC 0            REV 09   750-037355   CAAF0924         MPC Type 4-2
  CPU            REV 08   711-035209   CAAB9842         HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-021308   19T511101656     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AMA04RU          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   193363A00558     SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B10M00202        SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP
    Xcvr 0                NON-JNPR     X12J00328        CFP-100G-SR10
  PIC 2                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-031980   AMA088W          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B10L04211        SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   19T511101602     SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B10L04151        SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP
Xcvr 0               NON-JNPR     X12J00332        CFP-100G-SR10
FPC 1            REV 18   750-033205   ZE0128           MPC Type 3
  CPU            REV 06   711-035209   ZG5431           HMPC PMB 2G
  MIC 0          REV 15   750-033199   ZP6435           1X100GE CFP
    PIC 0                 BUILTIN      BUILTIN          1X100GE CFP
      Xcvr 0     REV 01   740-032210   J11E46118        CFP-100G-LR4
  MIC 1          REV 15   750-033199   ZP6442           1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN          1X100GE CFP
      Xcvr 0     REV 01   740-032210   UMN03T4          CFP-100G-LR4
FPC 2            REV 16   750-037358   CAAL1001         MPC Type 4-1
  CPU            REV 08   711-035209   CAAK7927         HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN          8X10GE SFPP
    Xcvr 0       REV 01   740-031980   193363A00589     SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   973152A00028     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   193363A00376     SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   973152A00016     SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   193363A00499     SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   973152A00039     SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   B11E01239        SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   973152A00058     SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          8X10GE SFPP
    Xcvr 0       REV 01   740-031980   B10M00075        SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   973152A00014     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AMA0638          SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   973152A00063     SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   AMA0629          SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   973152A00053     SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   193363A00344     SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   973152A00046     SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN          8X10GE SFPP
    Xcvr 0       REV 01   740-031980   AMA062M          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   973152A00080     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   193363A00580     SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   973152A00064     SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   093363A01494     SFP+-10G-SR
```



```
      Xcvr 5       REV 01   740-021308   973152A00020      SFP+-10G-SR
      Xcvr 6       REV 01   740-031980   123363A00047      SFP+-10G-SR
      Xcvr 7       REV 01   740-021308   973152A00072      SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-021308   03DZ06A01033      SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   973152A00022      SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   03DZ06A01026      SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   973152A00013      SFP+-10G-SR
      Xcvr 4       REV 01   740-021308   03DZ06A01028      SFP+-10G-SR
      Xcvr 5       REV 01   740-021308   973152A00079      SFP+-10G-SR
      Xcvr 6       REV 01   740-021308   03DZ06A01018      SFP+-10G-SR
      Xcvr 7       REV 01   740-021308   973152A00025      SFP+-10G-SR
  FPC 3            REV 33   750-028467   CAAF5400          MPC 3D 16x 10GE
    CPU            REV 11   711-029089   CAAH7626          AMPC PMB
    PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   973152A00066      SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   973152A00021      SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   973152A00062      SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   973152A00027      SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   973152A00065      SFP+-10G-SR
   Xcvr 1       REV 01   740-021308   973152A00069      SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   973152A00026      SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   973152A00003      SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   973152A00035      SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   973152A00004      SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   973152A00049      SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   973152A00055      SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   973152A00010      SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   973152A00001      SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   973152A00073      SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   973152A00012      SFP+-10G-SR
  FPC 4            REV 21   750-033205   ZG5028            MPC Type 3
    CPU            REV 05   711-035209   YX3911            HMPC PMB 2G
    MIC 0          REV 03   750-036233   ZL2036            2X40GE QSFP
      PIC 0                 BUILTIN      BUILTIN           2X40GE QSFP
        Xcvr 0     REV 01   740-032986   QB220708          QSFP+-40G-SR4
        Xcvr 1     REV 01   740-032986   QB220735          QSFP+-40G-SR4
    MIC 1          REV 03   750-036233   ZL2028            2X40GE QSFP
      PIC 2                 BUILTIN      BUILTIN           2X40GE QSFP
        Xcvr 0     REV 01   740-032986   QB220727          QSFP+-40G-SR4
        Xcvr 1     REV 01   740-032986   QB220715          QSFP+-40G-SR4
  FPC 5            REV 11   750-037358   CAAE2196          MPC Type 4-1
    CPU            REV 08   711-035209   CAAD9074          HMPC PMB 2G
    PIC 0                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-031980   AMA062S           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AMA062P           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AMA052R           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AMA0632           SFP+-10G-SR
      Xcvr 4       REV 01   740-031980   193363A00564      SFP+-10G-SR
      Xcvr 5       REV 01   740-031980   193363A00229      SFP+-10G-SR
      Xcvr 6       REV 01   740-031980   193363A00363      SFP+-10G-SR
      Xcvr 7       REV 01   740-031980   193363A00278      SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-031980   AMA04CC           SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   AD0927A001W       SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AMA04N2           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AMA062U           SFP+-10G-SR
      Xcvr 4       REV 01   740-031980   193363A00491      SFP+-10G-SR
```

```
      Xcvr 5       REV 01   740-031980   183363A01511      SFP+-10G-SR
      Xcvr 6       REV 01   740-031980   193363A00565      SFP+-10G-SR
      Xcvr 7       REV 01   740-031980   193363A00405      SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-031980   AMA07QX           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AMA06MS           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   193363A00318      SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   193363A00402      SFP+-10G-SR
      Xcvr 4       REV 01   740-031980   193363A00174      SFP+-10G-SR
      Xcvr 5       REV 01   740-031980   193363A00388      SFP+-10G-SR
      Xcvr 6       REV 01   740-031980   193363A00377      SFP+-10G-SR
      Xcvr 7       REV 01   740-031980   193363A00234      SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-031980   AMA062T           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   193363A00550      SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   193363A00364      SFP+-10G-SR
  Xcvr 3       REV 01   740-031980   AMA0630           SFP+-10G-SR
      Xcvr 4       REV 01   740-031980   193363A00509      SFP+-10G-SR
      Xcvr 5       REV 01   740-031980   193363A00459      SFP+-10G-SR
      Xcvr 6       REV 01   740-031980   113363A00191      SFP+-10G-SR
      Xcvr 7       REV 01   740-031980   193363A00352      SFP+-10G-SR
  FPC 6            REV 33   750-028467   CAAF5552          MPC 3D 16x 10GE
    CPU            REV 11   711-029089   CAAH7601          AMPC PMB
    PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   AD0927A0036       SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   AD0927A003M       SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   AD0927A003G       SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   AD0927A0031       SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   193363A00331      SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   193363A00325      SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   193363A00417      SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   183363A02509      SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   T09K75140         SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11A04356         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K01952         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K01914         SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   T09K75157         SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   T09K75194         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K01926         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K01936         SFP+-10G-SR
  FPC 7            REV 16   750-037358   CAAL1012          MPC Type 4-1
    CPU            REV 08   711-035209   CAAJ3851          HMPC PMB 2G
    PIC 0                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-031980   AMA04NK           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11F00260         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11E02192         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AMA04CP           SFP+-10G-SR
      Xcvr 4       REV 01   740-031980   AJ40JKK           SFP+-10G-SR
      Xcvr 5       REV 01   740-031980   B11F00238         SFP+-10G-SR
      Xcvr 6       REV 01   740-031980   B10M00275         SFP+-10G-SR
      Xcvr 7       REV 01   740-031980   193363A00211      SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           8X10GE SFPP
      Xcvr 0       REV 01   740-031980   B11D05577         SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11G00586         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AMA08B7           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AMA04Q0           SFP+-10G-SR
      Xcvr 4       REV 01   740-031980   B11D05840         SFP+-10G-SR
      Xcvr 5       REV 01   740-031980   B11E00467         SFP+-10G-SR
```

```
    Xcvr 6       REV 01   740-031980   B11E00029        SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   19T511101712     SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN          8X10GE SFPP
    Xcvr 0       REV 01   740-031980   193363A00568     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B10M00166        SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B10M00212        SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11D05823        SFP+-10G-SR
    Xcvr 4       REV 01   740-021308   03DZ06A01005     SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   03DZ06A01003     SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   03DZ06A01009     SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   03DZ06A01004     SFP+-10G-SR
PIC 3                   BUILTIN      BUILTIN          8X10GE SFPP
    Xcvr 0       REV 01   740-021308   03DZ06A01017     SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   03DZ06A01016     SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   03DZ06A01024     SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   03DZ06A01008     SFP+-10G-SR
    Xcvr 4       REV 01   740-030658   AD0946A02UH      SFP+-10G-USR
    Xcvr 5       REV 01   740-021308   T09J67913        SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AD0837ES09G      SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   03DZ06A01015     SFP+-10G-SR
FPC 8            REV 03   750-045372   CAAD3111         MPC Type 3
  CPU            REV 08   711-035209   CAAD8033         HMPC PMB 2G
  MIC 0          REV 03   750-036233   ZL2032           2X40GE QSFP
    PIC 0                 BUILTIN      BUILTIN          2X40GE QSFP
      Xcvr 0     REV 01   740-032986   QB230273         QSFP+-40G-SR4
      Xcvr 1     REV 01   740-032986   QB230254         QSFP+-40G-SR4
  MIC 1          REV 03   750-036233   ZL2021           2X40GE QSFP
    PIC 2                 BUILTIN      BUILTIN          2X40GE QSFP
      Xcvr 0     REV 01   740-032986   QB390962         QSFP+-40G-SR4
      Xcvr 1     REV 01   740-032986   QB390960         QSFP+-40G-SR4
FPC 9            REV 09   750-037355   CAAF1531         MPC Type 4-2
  CPU            REV 08   711-035209   CAAB9927         HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-031980   193363A00525     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   193363A00504     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   193363A00368     SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AJ40JSS          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP
  PIC 2                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-031980   123363A00042     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B10M00023        SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJ802EM          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11E02348        SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP
ADC 0            REV 13   750-043596   ABBX5532         Adapter Card
ADC 1            REV 13   750-043596   ABBX5550         Adapter Card
ADC 2            REV 13   750-043596   ABBX5571         Adapter Card
ADC 3            REV 13   750-043596   ABBX5568         Adapter Card
ADC 4            REV 13   750-043596   ABBX5556         Adapter Card
ADC 5            REV 13   750-043596   ABBX5553         Adapter Card
ADC 6            REV 13   750-043596   ABBX5541         Adapter Card
ADC 7            REV 13   750-043596   ABBX5578         Adapter Card
ADC 8            REV 13   750-043596   ABBX5560         Adapter Card
ADC 9            REV 07   750-043596   ABBV7188         Adapter Card
Fan Tray 0       REV 03   760-046960   ACAY0127         172mm FanTray - 6 Fans
Fan Tray 1       REV 2A   760-046960   ACAY0068         172mm FanTray - 6 Fans
Fan Tray 2       REV 2A   760-046960   ACAY0072         172mm FanTray - 6 Fans
Fan Tray 3       REV 2A   760-046960   ACAY0070         172mm FanTray - 6 Fans
```

show chassis hardware extensive (MX2010 Router)

```
user@host > show chassis hardware extensive
Hardware inventory:
Item                Version  Part number  Serial number     Description
Chassis                                   JN11E233DAFK      MX2010
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
                                  S/N:               JN11E233DAFK
 Assembly ID:  0x0557            Assembly Version:  00.00
 Date:          00-00-0000        Assembly Flags:    0x00
 ID: MX2010
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 31 45 32 33 33 44 41 46 4b 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 26   750-044636   ABAB9357          Lower Backplane
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           750-044636        S/N:               ABAB9357
 Assembly ID:  0x0b66            Assembly Version:  01.26
 Date:          08-28-2012        Assembly Flags:    0x00
 Version:       REV 26            CLEI Code:         PROTOXCLEI
 ID: Lower Backplane              FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ad 01 08 00 2c 21 72 70 a0 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 66 01 1a 52 45 56 20 32 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 35 37 00 1c 08 07
  Address 0x30: dc ff ff ff ad 01 08 00 2c 21 72 70 a0 00 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1       REV 01   711-044557   ABAB8643          Upper Backplane
 Jedec Code:    0x7fb0            EEPROM Version:    0x01
 P/N:           711-044557        S/N:               ABAB8643
 Assembly ID:  0x0b65            Assembly Version:  01.01
 Date:          07-27-2012        Assembly Flags:    0x00
 Version:       REV 01
 ID: Upper Backplane
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b 65 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 38 36 34 33 00 1b 07 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP              REV 04   711-032426   ACAJ1677          Power Midplane
Jedec Code:    0x7fb0            EEPROM Version:    0x01
 P/N:           711-032426        S/N:               ACAJ1677
```

```
 Assembly ID:  0x045d           Assembly Version:  01.04
 Date:         07-20-2012       Assembly Flags:    0x00
 Version:      REV 04
 ID: Power Midplane
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 4a 31 36 37 37 00 14 07 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 08   760-044634   ABBV9726          Front Panel Display
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          760-044634       S/N:               ABBV9726
 Assembly ID:  0x0b64           Assembly Version:  01.08
 Date:         09-10-2012       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         IPMYA4EJRA
 ID: Front Panel Display        FRU Model Number:  MX2010-CRAFT-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 64 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 56 39 37 32 36 00 0a 09 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
  Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0            REV 01   740-045050   1E02224000P       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-045050       S/N:               1E02224000P
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         12-06-2012       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         XXXXXXXXXX
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-HC-DC-S-A
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
  Address 0x20: 31 45 30 32 32 32 34 30 30 30 50 00 00 06 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
  Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1            REV 01   740-045050   1E02224000M       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-045050       S/N:               1E02224000M
Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         12-06-2012       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         XXXXXXXXXX
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-HC-DC-S-A
 Board Information Record:
```



```
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
  Address 0x20: 31 45 30 32 32 32 34 30 30 30 4d 00 00 06 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
  Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
...
PDM 0            REV 01   740-045234   1E262250067       DC Power Dist Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-045234       S/N:               1E262250067
 Assembly ID:  0x047b           Assembly Version:  01.01
 Date:         06-28-2012       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAJSKAA
 ID: DC Power Dist Module       FRU Model Number:  MX2000-PDM-DC-S-A
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 7b 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
  Address 0x20: 31 45 32 36 32 32 35 30 30 36 37 00 00 1c 06 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 2d 41
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 89 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 02   740-041821   9009099704        RE-S-1800x4
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-041821       S/N:               9009099704
 Assembly ID:  0x09c0           Assembly Version:  01.02
 Date:         03-15-2012       Assembly Flags:    0x00
 Version:      REV 02
 ID: RE-S-1800x4                FRU Model Number:  RE-S-1800X4-16G-S
 Board Information Record:
  Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
  Address 0x20: 39 30 30 39 30 39 39 37 30 34 00 00 00 0f 03 07
  Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
  Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3831 MB  UGB30SFA4000T1       SFA4000T1 00000651 Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI   11000019592       Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009099706        RE-S-1800x4
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-041821       S/N:               9009099706
 Assembly ID:  0x09c0           Assembly Version:  01.02
 Date:         02-23-2012       Assembly Flags:    0x00
 Version:      REV 02
```



```
 ID: RE-S-1800x4                FRU Model Number:  RE-S-1800X4-16G-S
 Board Information Record:
  Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
 Address 0x20: 39 30 30 39 30 39 39 37 30 36 00 00 00 17 02 07
  Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
  Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-499551-000404 Disk 1
CB 0             REV 13   750-040257   CAAF8436          Control Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-040257        S/N:               CAAF8436
 Assembly ID:  0x0b26            Assembly Version:  01.13
 Date:         08-29-2012        Assembly Flags:    0x00
 Version:      REV 13            CLEI Code:         PROTOXCLEI
 ID: Control Board               FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 26 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 35 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 46 38 34 33 36 00 1d 08 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
...
SPMB 0           REV 02   711-041855   ABBV3825          PMB Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          711-041855        S/N:               ABBV3825
 Assembly ID:  0x0b29            Assembly Version:  01.02
 Date:         08-14-2012        Assembly Flags:    0x00
 Version:      REV 02
 ID: PMB Board
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b 29 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 31 38 35 35 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 56 33 38 32 35 00 0e 08 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
...
SFB 0            REV 05   711-044466   ABBX5682          Switch Fabric Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          711-044466        S/N:               ABBX5682
 Assembly ID:  0x0b25            Assembly Version:  01.05
 Date:         09-07-2012        Assembly Flags:    0x00
 Version:      REV 05            CLEI Code:         PROTOXCLEI
 ID: Switch Fabric Board         FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```



```
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 25 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 34 36 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 58 35 36 38 32 00 07 09 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 00 00 00 01 00 00 00 00 00 00 48 00
...
FPC 0            REV 09   750-037355   CAAF0924          MPC Type 4-2
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-037355       S/N:               CAAF0924
 Assembly ID:  0x0b4e           Assembly Version:  01.09
 Date:         05-21-2012       Assembly Flags:    0x00
 Version:      REV 09           CLEI Code:         PROTOXCLEI
 ID: MPC Type 4-2               FRU Model Number:  MPC4E-2CGE-8XGE
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 4e 01 09 52 45 56 20 30 39 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 33 35 35 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 46 30 39 32 34 00 15 05 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
  Address 0x50: 50 43 34 45 2d 32 43 47 45 2d 38 58 47 45 00 00
  Address 0x60: 00 00 00 00 00 00 30 39 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c6 ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 08   711-035209   CAAB9842          HMPC PMB 2G
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          711-035209       S/N:               CAAB9842
 Assembly ID:  0x0b04           Assembly Version:  01.08
 Date:         05-17-2012       Assembly Flags:    0x00
 Version:      REV 08
 ID: HMPC PMB 2G
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 42 39 38 34 32 00 11 05 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  PIC 0                   BUILTIN      BUILTIN           4x10GE SFPP
 Jedec Code:   0x0000           EEPROM Version:    0x00
 P/N:          BUILTIN          S/N:               BUILTIN
 Assembly ID:  0x0a53           Assembly Version:  00.00
 Date:         00-00-0000       Assembly Flags:    0x00
 ID: 4x10GE SFPP
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 53 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
  Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```



```
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ae 64 00 00 00 00 0a 52 00 00
    Xcvr 0       REV 01   740-021308   19T511101656      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AMA04RU           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   193363A00558      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B10M00202         SFP+-10G-SR
...
ADC 0            REV 13   750-043596   ABBX5532          Adapter Card
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-043596        S/N:               ABBX5532
 Assembly ID:  0x0b3d            Assembly Version:  01.13
 Date:         09-12-2012        Assembly Flags:    0x00
 Version:      REV 13            CLEI Code:         IPUCBA8CAA
 ID: Adapter Card                FRU Model Number:  MX2000-LC-ADAPTER
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 3d 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 58 35 35 33 32 00 0c 09 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00 00
...
```

show chassis hardware models (MX2010 Router)

```
user@host > show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
FPM Board        REV 06   711-032349   ZX8744            711-032349
PSM 4            REV 0C   740-033727   VK00254           00000000000000000000000
PSM 5            REV 0B   740-033727   VG00015           00000000000000000000000
PSM 6            REV 0B   740-033727   VH00097           00000000000000000000000
PSM 7            REV 0C   740-033727   VJ00151           00000000000000000000000
PSM 8            REV 0C   740-033727   VJ00149           00000000000000000000000
PDM 0            REV 0B   740-038109   WA00008
PDM 1            REV 0B   740-038109   WA00014
Routing Engine 0 REV 02   740-041821   9009094134        RE-S-1800X4-16G-S
Routing Engine 1 REV 02   740-041821   9009094141        RE-S-1800X4-16G-S
CB 0             REV 08   750-040257   CAAB3491          750-040257
CB 1             REV 08   750-040257   CAAB3489          750-040257
SFB 0            REV 06   711-032385   ZV1828            711-032385
SFB 1            REV 07   711-032385   ZZ2568            711-032385
SFB 2            REV 07   711-032385   ZZ2563            711-032385
SFB 3            REV 07   711-032385   ZZ2564            711-032385
SFB 4            REV 07   711-032385   ZZ2580            711-032385
SFB 5            REV 07   711-032385   ZZ2579            711-0323856
SFB 6            REV 07   711-032385   CAAB4882          711-044170
SFB 7            REV 07   711-032385   CAAB4898          711-044170
FPC 0            REV 33   750-028467   CAAB1919          MPC-3D-16XGE-SFPP
FPC 1            REV 21   750-033205   ZG5027            MX-MPC3-3D
  MIC 0          REV 03   750-033307   ZV6299            MIC3-3D-10XGE-SFPP
  MIC 1          REV 03   750-033307   ZV6268            MIC3-3D-10XGE-SFPP
FPC 8            REV 22   750-031089   ZT9746            MX-MPC2-3D
  MIC 0          REV 26   750-028392   ABBS1150          MIC-3D-20GE-SFP
  MIC 1          REV 26   750-028387   ABBR9582          MIC-3D-4XGE-XFP
FPC 9            REV 11   750-036284   ZL3591            MPCE-3D-16XGE-SFPP
ADC 0            REV 05   750-043596   CAAC2073          750-043596
```



```
ADC 1            REV 01   750-043596   ZV4117            750-043596
ADC 8            REV 01   750-043596   ZV4107            750-043596
ADC 9            REV 02   750-043596   ZW1555            750-043596
Fan Tray 0       REV 2A   760-046960   ACAY0015
Fan Tray 1       REV 2A   760-046960   ACAY0019
Fan Tray 2       REV 2A   760-046960   ACAY0020
Fan Tray 3       REV 2A   760-046960   ACAY0021
```

show chassis hardware clei-models (MX2010 Routers)

```
user@host > show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
FPM Board        REV 06   711-032349   PROTOXCLEI        711-032349
PSM 4            REV 0C   740-033727   0000000000        0000000000000000000000
PSM 5            REV 0B   740-033727   0000000000        0000000000000000000000
PSM 6            REV 0B   740-033727   0000000000        0000000000000000000000
PSM 7            REV 0C   740-033727   0000000000        0000000000000000000000
PSM 8            REV 0C   740-033727   0000000000        0000000000000000000000
PDM 0            REV 0B   740-038109
PDM 1            REV 0B   740-038109
Routing Engine 0 REV 02   740-041821                     RE-S-1800X4-16G-S
Routing Engine 1 REV 02   740-041821                     RE-S-1800X4-16G-S
CB 0             REV 08   750-040257   PROTOXCLEI        750-040257
CB 1             REV 08   750-040257   PROTOXCLEI        750-040257
SFB 0            REV 06   711-032385   PROTOXCLEI        711-032385
SFB 1            REV 07   711-032385   PROTOXCLEI        711-032385
SFB 2            REV 07   711-032385   PROTOXCLEI        711-032385
SFB 3            REV 07   711-032385   PROTOXCLEI        711-032385
SFB 4            REV 07   711-032385   PROTOXCLEI        711-032385
SFB 5            REV 07   711-032385   PROTOXCLEI        711-0323856
SFB 6            REV 07   711-032385   PROTOXCLEI        711-044170
SFB 7            REV 07   711-032385   PROTOXCLEI        711-044170
FPC 0            REV 33   750-028467                     MPC-3D-16XGE-SFPP
FPC 1            REV 21   750-033205                     MX-MPC3-3D
  MIC 0          REV 03   750-033307   PROTOXCLEI        MIC3-3D-10XGE-SFPP
  MIC 1          REV 03   750-033307   PROTOXCLEI        MIC3-3D-10XGE-SFPP
FPC 8            REV 22   750-031089   COUIBAYBAA        MX-MPC2-3D
  MIC 0          REV 26   750-028392   COUIA15BAA        MIC-3D-20GE-SFP
  MIC 1          REV 26   750-028387   COUIA16BAA        MIC-3D-4XGE-XFP
FPC 9            REV 11   750-036284   CMUIACGBAA        MPCE-3D-16XGE-SFPP
ADC 0            REV 05   750-043596   PROTOXCLEI        750-043596
ADC 1            REV 01   750-043596   PROTOXCLEI        750-043596
ADC 8            REV 01   750-043596   PROTOXCLEI        750-043596
ADC 9            REV 02   750-043596   PROTOXCLEI        750-043596
Fan Tray 0       REV 2A   760-046960
Fan Tray 1       REV 2A   760-046960
Fan Tray 2       REV 2A   760-046960
Fan Tray 3       REV 2A   760-046960
```

show chassis hardware (MX2010 Routers with MPC6E and OTN MIC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11C9AFEAFK      MX2010
Midplane         REV 35   750-044636   ABAB9188          Lower Backplane
Midplane 1       REV 02   711-044557   ABAB8729          Upper Backplane
PMP              REV 04   711-032426   ACAJ2432          Power Midplane
FPM Board        REV 09   760-044634   ABCA4314          Front Panel Display
```



```
PSM 0            REV 01   740-050037   1EDB321015C       DC 52V Power Supply
Module
PSM 1            REV 01   740-050037   1EDB321015J       DC 52V Power Supply
Module
PSM 2            REV 01   740-050037   1EDB32000K8       DC 52V Power Supply
Module
PSM 3            REV 01   740-050037   1EDB32101JW       DC 52V Power Supply
Module
PSM 4            REV 01   740-050037   1EDB321015G       DC 52V Power Supply
Module
PSM 5            REV 01   740-050037   1EDB32101HH       DC 52V Power Supply
Module
PSM 6            REV 01   740-050037   1EDB32101HD       DC 52V Power Supply
Module
PSM 7            REV 01   740-050037   1EDB321015F       DC 52V Power Supply
Module
PSM 8            REV 01   740-050037   1EDB321015B       DC 52V Power Supply
Module
PDM 0            REV 03   740-045234   1EFA3220433       DC Power Dist Module
PDM 1            REV 03   740-045234   1EFA3220425       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009115685        RE-S-1800x4
Routing Engine 1 REV 02   740-041821   9009099711        RE-S-1800x4
CB 0             REV 23   750-040257   CABE8395          Control Board
CB 1             REV 12   750-040257   CAAD9499          Control Board
SPMB 0           REV 02   711-041855   ABCG8426          PMB Board
SPMB 1           REV 02   711-041855   ABBS1481          PMB Board
SFB 0            REV 06   711-044466   ABCD5013          Switch Fabric Board
SFB 1            REV 06   711-044466   ABCD5160          Switch Fabric Board
SFB 2            REV 06   711-044466   ABCD5175          Switch Fabric Board
SFB 3            REV 06   711-044466   ABCD4938          Switch Fabric Board
SFB 4            REV 06   711-044466   ABCD4944          Switch Fabric Board
SFB 5            REV 06   711-044466   ABCD4968          Switch Fabric Board
SFB 6            REV 06   711-044466   ABCD5267          Switch Fabric Board
SFB 7            REV 06   711-044466   ABCD4997          Switch Fabric Board
FPC 0            REV 59   750-044130   ABCT7676          MPC6E 3D
  CPU            REV 10   711-045719   ABCK8527          RMPC PMB
  XLM 0          REV 13   711-046638   ABCT7810          MPC6E XL
  XLM 1          REV 13   711-046638   ABCT7811          MPC6E XL
FPC 2            REV 27   750-033205   ZL6014            MPCE Type 3 3D
  CPU            REV 07   711-035209   ZK9068            HMPC PMB 2G
  MIC 0          REV 14   750-033196   CAAW9214          1X100GE CXP
    PIC 0                 BUILTIN      BUILTIN           1X100GE CXP
      Xcvr 0     REV 01   740-046563   XC49FC030         CFP2-100G-SR10
  MIC 1          REV 18   750-033199   CAAC3231          1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN           1X100GE CFP
FPC 3            REV 59   750-044130   ABCT7682          MPC6E 3D
  CPU            REV 10   711-045719   ABCK8531          RMPC PMB
  XLM 0          REV 13   711-046638   ABCT7818          MPC6E XL
  XLM 1          REV 13   711-046638   ABCT7819          MPC6E XL
FPC 4            REV 33   750-044130   ABBY9278          MPC6E 3D
  CPU            REV 09   711-045719   ABBY8677          RMPC PMB
  XLM 0          REV 06.2.00 711-046638 ABBY8844         MPC6E XL
  XLM 1          REV 06.2.00 711-046638 ABBY8830         MPC6E XL
FPC 5            REV 59   750-044130   ABCT7675          MPC6E 3D
  CPU            REV 10   711-045719   ABCK8526          RMPC PMB
  XLM 0          REV 13   711-046638   ABCT7808          MPC6E XL
  XLM 1          REV 13   711-046638   ABCT7809          MPC6E XL
FPC 6            REV 30   750-028467   ZM4986            MPC 3D 16x 10GE
 CPU            REV 10   711-029089   ZP6541            AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AQ43GAC           SFP+-10G-SR
```

```
  PIC 1                        BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
    Xcvr 0        REV 01       740-031980   ALM0A6D            SFP+-10G-SR
    Xcvr 1        REV 01       740-031980   AQF0RB3            SFP+-10G-SR
    Xcvr 2        REV 01       740-031980   153363A00333       SFP+-10G-SR
    Xcvr 3        REV 01       740-021308   AN10KYE            SFP+-10G-SR
  PIC 2                        BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
    Xcvr 0        REV 01       740-021308   APK04YM            SFP+-10G-SR
  PIC 3                        BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
    Xcvr 0        REV 01       740-031980   AQF0H44            SFP+-10G-SR
FPC 8             REV 38       750-031090   CABF7313           MPC Type 2 3D EQ
  CPU             REV 08       711-030884   CABE6727           MPC PMB 2G
  MIC 0           REV 18       750-028380   YK8253             3D 2x 10GE XFP
    PIC 0                      BUILTIN      BUILTIN            1x 10GE XFP
      Xcvr 0      REV 03       740-014289   AD1148M00TP        XFP-10G-SR
    PIC 1                      BUILTIN      BUILTIN            1x 10GE XFP
  QXM 0           REV 06       711-028408   CABC5614           MPC QXM
  QXM 1           REV 06       711-028408   CABC5550           MPC QXM
FPC 9             REV 39       750-044130   ABCK1652           MPC6E 3D
  CPU             REV 09       711-045719   ABCK1655           RMPC PMB
  MIC 0           REV 09       750-049457   ABCP1230           2X100GE CFP2 OTN
    PIC 0                      BUILTIN      BUILTIN            2X100GE CFP2 OTN
      Xcvr 0                   NON-JNPR     37300222WP0002     CFP2-100G-LR4-D
      Xcvr 1                   NON-JNPR     FD46F001Y          CFP2-100G-SR10
  MIC 1           REV 07       750-049457   ABCV6662           2X100GE CFP2 OTN
    PIC 1                      BUILTIN      BUILTIN            2X100GE CFP2 OTN
      Xcvr 0                   NON-JNPR     UQD0014            CFP2-100G-LR4-D
      Xcvr 1                   NON-JNPR     J13J68335          CFP2-100G-LR4-D
  XLM 0           REV 07.2.00 711-046638 ABCK5491              MPC6E XL
  XLM 1           REV 07.2.00 711-046638 ABCK5475              MPC6E XL
ADC 1             REV 17       750-043596   ABCG9023           Adapter Card
ADC 2             REV 01       750-043596   ZV4079             Adapter Card
ADC 6             REV 17       750-043596   ABCG8866           Adapter Card
ADC 8             REV 17       750-043596   ABCA8993           Adapter Card
Fan Tray 0        REV 06       760-046960   ACAY0354           172mm FanTray - 6 Fans
Fan Tray 1        REV 06       760-046960   ACAY0831           172mm FanTray - 6 Fans
Fan Tray 2        REV 06       760-046960   ACAY0892           172mm FanTray - 6 Fans
Fan Tray 3        REV 06       760-046960   ACAY0839           172mm FanTray - 6 Fans
```

show chassis hardware detail (MX2010 Routers with MPC6E and OTN MIC)

```
user@host> show chassis hardware detail
Hardware inventory:
Item              Version      Part number  Serial number      Description
Chassis                                     JN11C9AFEAFK       MX2010
Midplane          REV 35       750-044636   ABAB9188           Lower Backplane
Midplane 1        REV 02       711-044557   ABAB8729           Upper Backplane
PMP               REV 04       711-032426   ACAJ2432           Power Midplane
FPM Board         REV 09       760-044634   ABCA4314           Front Panel Display
PSM 0             REV 01       740-050037   1EDB321015C        DC 52V Power Supply
Module
PSM 1             REV 01       740-050037   1EDB321015J        DC 52V Power Supply
Module
PSM 2             REV 01       740-050037   1EDB32000K8        DC 52V Power Supply
Module
PSM 3             REV 01       740-050037   1EDB32101JW        DC 52V Power Supply
Module
PSM 4             REV 01       740-050037   1EDB321015G        DC 52V Power Supply
Module
PSM 5             REV 01       740-050037   1EDB32101HH        DC 52V Power Supply
Module
PSM 6             REV 01       740-050037   1EDB32101HD        DC 52V Power Supply
```

```
Module
PSM 7            REV 01   740-050037   1EDB321015F       DC 52V Power Supply
Module
PSM 8            REV 01   740-050037   1EDB321015B       DC 52V Power Supply
Module
PDM 0            REV 03   740-045234   1EFA3220433       DC Power Dist Module
PDM 1            REV 03   740-045234   1EFA3220425       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009115685        RE-S-1800x4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200274310822 191 Compact Flash
  ad1   30533 MB  UGB94BPH32H0S1-KCI   11000043190       Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 1         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 2         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
  DIMM 3         VL31B5263F-F8SD DIE REV-0 PCB REV-0     MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009099711        RE-S-1800x4
  ad0    3998 MB  Virtium - TuffDrive VCF P1T0200262860208 30 Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-499551-000146 Disk 1
CB 0             REV 23   750-040257   CABE8395          Control Board
CB 1             REV 12   750-040257   CAAD9499          Control Board
SPMB 0           REV 02   711-041855   ABCG8426          PMB Board
SPMB 1           REV 02   711-041855   ABBS1481          PMB Board
SFB 0            REV 06   711-044466   ABCD5013          Switch Fabric Board
SFB 1            REV 06   711-044466   ABCD5160          Switch Fabric Board
SFB 2            REV 06   711-044466   ABCD5175          Switch Fabric Board
SFB 3            REV 06   711-044466   ABCD4938          Switch Fabric Board
SFB 4            REV 06   711-044466   ABCD4944          Switch Fabric Board
SFB 5            REV 06   711-044466   ABCD4968          Switch Fabric Board
SFB 6            REV 06   711-044466   ABCD5267          Switch Fabric Board
SFB 7            REV 06   711-044466   ABCD4997          Switch Fabric Board
FPC 0            REV 59   750-044130   ABCT7676          MPC6E 3D
  CPU            REV 10   711-045719   ABCK8527          RMPC PMB
  XLM 0          REV 13   711-046638   ABCT7810          MPC6E XL
  XLM 1          REV 13   711-046638   ABCT7811          MPC6E XL
FPC 2            REV 27   750-033205   ZL6014            MPCE Type 3 3D
  CPU            REV 07   711-035209   ZK9068            HMPC PMB 2G
  MIC 0          REV 14   750-033196   CAAW9214          1X100GE CXP
    PIC 0                 BUILTIN      BUILTIN           1X100GE CXP
      Xcvr 0     REV 01   740-046563   XC49FC030         CFP2-100G-SR10
  MIC 1          REV 18   750-033199   CAAC3231          1X100GE CFP
    PIC 2                 BUILTIN      BUILTIN           1X100GE CFP
FPC 3            REV 59   750-044130   ABCT7682          MPC6E 3D
  CPU            REV 10   711-045719   ABCK8531          RMPC PMB
  XLM 0          REV 13   711-046638   ABCT7818          MPC6E XL
 XLM 1          REV 13   711-046638   ABCT7819          MPC6E XL
FPC 4            REV 33   750-044130   ABBY9278          MPC6E 3D
  CPU            REV 09   711-045719   ABBY8677          RMPC PMB
  XLM 0          REV 06.2.00 711-046638 ABBY8844         MPC6E XL
  XLM 1          REV 06.2.00 711-046638 ABBY8830         MPC6E XL
FPC 5            REV 59   750-044130   ABCT7675          MPC6E 3D
  CPU            REV 10   711-045719   ABCK8526          RMPC PMB
  XLM 0          REV 13   711-046638   ABCT7808          MPC6E XL
  XLM 1          REV 13   711-046638   ABCT7809          MPC6E XL
FPC 6            REV 30   750-028467   ZM4986            MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ZP6541            AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AQ43GAC           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   ALM0A6D           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AQF0RB3           SFP+-10G-SR
```

```
    Xcvr 2         REV 01   740-031980   153363A00333      SFP+-10G-SR
    Xcvr 3         REV 01   740-021308   AN10KYE           SFP+-10G-SR
  PIC 2                     BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0         REV 01   740-021308   APK04YM           SFP+-10G-SR
  PIC 3                     BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0         REV 01   740-031980   AQF0H44           SFP+-10G-SR
FPC 8              REV 38   750-031090   CABF7313          MPC Type 2 3D EQ
  CPU              REV 08   711-030884   CABE6727          MPC PMB 2G
  MIC 0            REV 18   750-028380   YK8253            3D 2x 10GE XFP
    PIC 0                   BUILTIN      BUILTIN           1x 10GE XFP
      Xcvr 0       REV 03   740-014289   AD1148M00TP       XFP-10G-SR
    PIC 1                   BUILTIN      BUILTIN           1x 10GE XFP
  QXM 0            REV 06   711-028408   CABC5614          MPC QXM
  QXM 1            REV 06   711-028408   CABC5550          MPC QXM
FPC 9              REV 39   750-044130   ABCK1652          MPC6E 3D
  CPU              REV 09   711-045719   ABCK1655          RMPC PMB
  MIC 0            REV 09   750-049457   ABCP1230          2X100GE CFP2 OTN
    PIC 0                   BUILTIN      BUILTIN           2X100GE CFP2 OTN
      Xcvr 0                NON-JNPR     37300222WP0002    CFP2-100G-LR4-D
      Xcvr 1                NON-JNPR     FD46F001Y         CFP2-100G-SR10
  MIC 1            REV 07   750-049457   ABCV6662          2X100GE CFP2 OTN
    PIC 1                   BUILTIN      BUILTIN           2X100GE CFP2 OTN
      Xcvr 0                NON-JNPR     UQD0014           CFP2-100G-LR4-D
      Xcvr 1                NON-JNPR     J13J68335         CFP2-100G-LR4-D
  XLM 0            REV 07.2.00 711-046638 ABCK5491         MPC6E XL
  XLM 1            REV 07.2.00 711-046638 ABCK5475         MPC6E XL
ADC 1              REV 17   750-043596   ABCG9023          Adapter Card
ADC 2              REV 01   750-043596   ZV4079            Adapter Card
ADC 6              REV 17   750-043596   ABCG8866          Adapter Card
ADC 8              REV 17   750-043596   ABCA8993          Adapter Card
Fan Tray 0         REV 06   760-046960   ACAY0354          172mm FanTray - 6 Fans
Fan Tray 1         REV 06   760-046960   ACAY0831          172mm FanTray - 6 Fans
Fan Tray 2         REV 06   760-046960   ACAY0892          172mm FanTray - 6 Fans
Fan Tray 3         REV 06   760-046960   ACAY0839          172mm FanTray - 6 Fans
```

show chassis hardware extensive (MX2010 Routers with MPC6E and OTN MIC)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item               Version  Part number  Serial number     Description
Chassis                                  JN11C9AFEAFK      MX2010
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
                                  S/N:               JN11C9AFEAFK
 Assembly ID:   0x0557            Assembly Version:  00.00
 Date:          00-00-0000        Assembly Flags:    0x00
 ID: MX2010
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 31 43 39 41 46 45 41 46 4b 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane           REV 35   750-044636   ABAB9188          Lower Backplane
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
 P/N:           750-044636        S/N:               ABAB9188
 Assembly ID:   0x0b66            Assembly Version:  01.35
```



```
  Date:         06-21-2013        Assembly Flags:    0x00
  Version:      REV 35            CLEI Code:         IPMU810ARA
  ID: Lower Backplane             FRU Model Number:  CHAS-BP-MX2010-S
  Board Information Record:
   Address 0x00: ad 01 08 00 3c 8a b0 38 68 00 ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 66 01 23 52 45 56 20 33 35 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
   Address 0x20: 53 2f 4e 20 41 42 41 42 39 31 38 38 00 15 06 07
   Address 0x30: dd ff ff ff ad 01 08 00 3c 8a b0 38 68 00 ff ff
   Address 0x40: ff ff ff ff 01 49 50 4d 55 38 31 30 41 52 41 43
   Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 31 30 2d 53 00
   Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff f8 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1       REV 02   711-044557   ABAB8729          Upper Backplane
  Jedec Code:   0x7fb0            EEPROM Version:    0x01
  P/N:          711-044557        S/N:               ABAB8729
  Assembly ID:  0x0b65            Assembly Version:  01.02
  Date:         03-21-2013        Assembly Flags:    0x00
  Version:      REV 02
  ID: Upper Backplane
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 01 ff 0b 65 01 02 52 45 56 20 30 32 00 00
   Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
   Address 0x20: 53 2f 4e 20 41 42 41 42 38 37 32 39 00 15 03 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
   Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP              REV 04   711-032426   ACAJ2432          Power Midplane
  Jedec Code:   0x7fb0            EEPROM Version:    0x01
  P/N:          711-032426        S/N:               ACAJ2432
  Assembly ID:  0x045d            Assembly Version:  01.04
  Date:         03-28-2013        Assembly Flags:    0x00
  Version:      REV 04
  ID: Power Midplane
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
   Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
   Address 0x20: 53 2f 4e 20 41 43 41 4a 32 34 33 32 00 1c 03 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
   Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
   Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
   Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 09   760-044634   ABCA4314          Front Panel Display
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          760-044634        S/N:               ABCA4314
  Assembly ID:  0x0b64            Assembly Version:  01.09
  Date:         03-28-2013        Assembly Flags:    0x00
  Version:      REV 09            CLEI Code:         IPMYA4EJRA
  ID: Front Panel Display         FRU Model Number:  MX2010-CRAFT-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 64 01 09 52 45 56 20 30 39 00 00
```

```
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 43 41 34 33 31 34 00 1c 03 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
  Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0            REV 01   740-050037   1EDB321015C       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB321015C
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-28-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 35 43 00 00 1c 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1            REV 01   740-050037   1EDB321015J       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB321015J
Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-28-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 35 4a 00 00 1c 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2            REV 01   740-050037   1EDB32000K8       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB32000K8
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-23-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 30 30 30 4b 38 00 00 17 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
```



```
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3            REV 01   740-050037   1EDB32101JW       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB32101JW
 Assembly ID:  0x0478            Assembly Version:  01.01
 Date:         05-30-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 4a 57 00 00 1e 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 4            REV 01   740-050037   1EDB321015G       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB321015G
 Assembly ID:  0x0478            Assembly Version:  01.01
 Date:         05-28-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 35 47 00 00 1c 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 5            REV 01   740-050037   1EDB32101HH       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB32101HH
 Assembly ID:  0x0478            Assembly Version:  01.01
 Date:         05-30-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 48 48 00 00 1e 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 6            REV 01   740-050037   1EDB32101HD       DC 52V Power Supply
```

```
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB32101HD
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-30-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 48 44 00 00 1e 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7            REV 01   740-050037   1EDB321015F       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB321015F
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-28-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 35 46 00 00 1c 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 8            REV 01   740-050037   1EDB321015B       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB321015B
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-28-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 31 30 31 35 42 00 00 1c 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PDM 0            REV 03   740-045234   1EFA3220433       DC Power Dist Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-045234       S/N:               1EFA3220433
 Assembly ID:  0x047b           Assembly Version:  01.03
 Date:         05-30-2013       Assembly Flags:    0x00
```



```
 Version:      REV 03            CLEI Code:         IPUPAJSKAA
 ID: DC Power Dist Module        FRU Model Number:  MX2000-PDM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 7b 01 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
  Address 0x20: 31 45 46 41 33 32 32 30 34 33 33 00 00 1e 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 33 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 1d 00 00 00 00 00 00 00 00 00 00 00 00
PDM 1            REV 03   740-045234   1EFA3220425       DC Power Dist Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-045234        S/N:               1EFA3220425
 Assembly ID:  0x047b            Assembly Version:  01.03
 Date:         05-30-2013        Assembly Flags:    0x00
 Version:      REV 03            CLEI Code:         IPUPAJSKAA
 ID: DC Power Dist Module        FRU Model Number:  MX2000-PDM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
..
```

show chassis hardware (MX2020 Router)

```
user@host > show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11E2227AFJ      MX2020
Midplane         REV 27   750-040240   ABAB9384          Lower Power Midplane
Midplane 1       REV 04   711-032386   ABAB9386          Upper Backplane
PMP 1            REV 05   711-032428   ACAJ1579          Upper Power Midplane
PMP 0            REV 04   711-032426   ACAJ1524          Lower Power Midplane
FPM Board        REV 06   760-040242   ABBT8837          Front Panel Display
PSM 0            REV 01   740-045050   1E022240056       DC 52V Power Supply
Module
PSM 1            REV 01   740-045050   1E022240054       DC 52V Power Supply
Module
PSM 2            REV 01   740-045050   1E02224005H       DC 52V Power Supply
Module
PSM 3            REV 01   740-045050   1E022240053       DC 52V Power Supply
Module
PSM 4            REV 01   740-045050   1E02224004K       DC 52V Power Supply
Module
PSM 7            REV 01   740-045050   1E02224006W       DC 52V Power Supply
Module
PSM 8            REV 01   740-045050   1E022240062       DC 52V Power Supply
Module
PSM 9            REV 01   740-045050   1E02224005B       DC 52V Power Supply
Module
PSM 10           REV 01   740-045050   1E02224005A       DC 52V Power Supply
Module
PSM 11           REV 01   740-045050   1E022240052       DC 52V Power Supply
Module
PSM 12           REV 01   740-045050   1E022240051       DC 52V Power Supply
Module
PSM 13           REV 01   740-045050   1E022240058       DC 52V Power Supply
Module
PSM 14           REV 01   740-045050   1E02224004L       DC 52V Power Supply
Module
```

```
PSM 15           REV 01   740-045050   1E02224005M       DC 52V Power Supply
Module
PSM 16           REV 01   740-045050   1E02224006S       DC 52V Power Supply
Module
PSM 17           REV 01   740-045050   1E02224005Z       DC 52V Power Supply
Module
PDM 0            REV 01   740-045234   1E012150033       DC Power Dist Module
PDM 1            REV 01   740-045234   1E012150027       DC Power Dist Module
PDM 2            REV 01   740-045234   1E012150028       DC Power Dist Module
PDM 3            REV 01   740-045234   1E012150045       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009089704        RE-S-1800x4
Routing Engine 1 REV 02   740-041821   9009094138        RE-S-1800x4
CB 0             REV 14   750-040257   CAAF8430          Control Board
CB 1             REV 08   750-040257   CAAB3482          Control Board
SPMB 0           REV 01   711-041855   ZS2290            PMB Board
SPMB 1           REV 02   711-041855   CAAA6141          PMB Board
SFB 0            REV 03   711-044466   ABBV6789          Switch Fabric Board
SFB 1            REV 05   711-044466   ABBX5666          Switch Fabric Board
SFB 2            REV 05   711-044466   ABBX5678          Switch Fabric Board
SFB 3            REV 05   711-044466   ABBX5687          Switch Fabric Board
SFB 4            REV 05   711-044466   ABBX5609          Switch Fabric Board
SFB 5            REV 05   711-044466   ABBX5675          Switch Fabric Board
SFB 6            REV 03   711-044466   ABBV6805          Switch Fabric Board
SFB 7            REV 05   711-044466   ABBX5701          Switch Fabric Board
FPC 0            REV 30   750-028467   ABBN0284          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBN0507          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E00990         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E04357         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01327         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E04375         SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E02760         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E02904         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E03963         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E00756         SFP+-10G-USR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E04418         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E01077         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E01128         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01253         SFP+-10G-USR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E01140         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11F01626         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E01075         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E01177         SFP+-10G-USR
FPC 1            REV 30   750-028467   ABBN0208          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBJ1084          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E04745         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11F01570         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E04388         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01439         SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E04739         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11F01869         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01675         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01901         SFP+-10G-USR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11F01346         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11F01288         SFP+-10G-USR
```



```
    Xcvr 2       REV 01   740-030658   B11F01824         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E04312         SFP+-10G-USR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E02811         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E03847         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01495         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01265         SFP+-10G-USR
FPC 2            REV 30   750-028467   ZM5111            MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ZP6607            AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80LJA           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80MFZ           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80NKL           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80KF4           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80FBJ           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80MM2           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80LJV           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80NXV           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80N1H           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80NLS           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80FL5           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80NL9           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80NG2           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80KDU           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80MG1           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80MM0           SFP+-10G-SR
FPC 3            REV 30   750-028467   ABBN0302          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBN0495          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11F01581         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E01176         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01251         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E02752         SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E00786         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E01020         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E01023         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E02819         SFP+-10G-USR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E02812         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11D04437         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01279         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01333         SFP+-10G-USR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E00978         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E01018         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01784         SFP+-10G-USR
    Xcvr 3       REV 01   740-031980   AK80NKP           SFP+-10G-SR
FPC 4            REV 30   750-028467   ABBN0308          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBJ1095          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E04305         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E01147         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E01195         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01743         SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11F01892         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E02880         SFP+-10G-USR
```

```
      Xcvr 2       REV 01   740-030658   B11E00725          SFP+-10G-USR
      Xcvr 3       REV 01   740-030658   B11E01057          SFP+-10G-USR
    PIC 2                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-030658   B11E02816          SFP+-10G-USR
      Xcvr 1       REV 01   740-030658   B11C04501          SFP+-10G-USR
      Xcvr 2       REV 01   740-030658   B11E02764          SFP+-10G-USR
      Xcvr 3       REV 01   740-030658   B11E00789          SFP+-10G-USR
    PIC 3                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-030658   B11F01250          SFP+-10G-USR
      Xcvr 1       REV 01   740-030658   B11E02847          SFP+-10G-USR
      Xcvr 2       REV 01   740-030658   B11E00787          SFP+-10G-USR
      Xcvr 3       REV 01   740-030658   B11E03803          SFP+-10G-USR
  FPC 5            REV 30   750-028467   ABBN0316           MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBJ1082           AMPC PMB
    PIC 0                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   B11K00523          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11K01848          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K01865          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K00540          SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   B11K00422          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11K00428          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K00423          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K01855          SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   B11K01847          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11K00526          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K00529          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K00525          SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   B11K00425          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11K00530          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K01851          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K00528          SFP+-10G-SR
  FPC 6            REV 32   750-028467   ABBN6832           MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBK6534           AMPC PMB
    PIC 0                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80MB4            SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80FQ6            SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80N1F            SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80NLQ            SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80KDR            SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80FGJ            SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80N5G            SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80KD8            SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80LET            SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80N1X            SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NRF            SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80NL2            SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80N3D            SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80MRB            SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80LEQ            SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80LER            SFP+-10G-SR
  FPC 7            REV 32   750-028467   ABBN6811           MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBN7288           AMPC PMB
    PIC 0                   BUILTIN      BUILTIN            4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80NK8            SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80LJG            SFP+-10G-SR
```



```
      Xcvr 2      REV 01   740-031980   AK80LBU          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80N21          SFP+-10G-SR
    PIC 1                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80LEU          SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80NLM          SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80NL6          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80LES          SFP+-10G-SR
    PIC 2                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80LEN          SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80ME0          SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80LMG          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80MM1          SFP+-10G-SR
    PIC 3                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80MG7          SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80KF9          SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80NRQ          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80NLE          SFP+-10G-SR
  FPC 8           REV 23   750-028467   YN2977           MPC 3D 16x 10GE
    CPU           REV 10   711-029089   YP1856           AMPC PMB
    PIC 0                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   183363A00875     SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   183363A00851     SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   183363A00772     SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   183363A00882     SFP+-10G-SR
    PIC 1                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   183363A00735     SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   183363A00169     SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   183363A00726     SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   183363A00077     SFP+-10G-SR
    PIC 2                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   183363A00168     SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   183363A00676     SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   183363A00732     SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   183363A00091     SFP+-10G-SR
    PIC 3                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   183363A00725     SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   183363A00642     SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   183363A00871     SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   183363A00853     SFP+-10G-SR
  FPC 9           REV 32   750-028467   ABBN6798         MPC 3D 16x 10GE
    CPU           REV 10   711-029089   ABBK6556         AMPC PMB
    PIC 0                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-021308   9ZDZ06A00055     SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   183363A00239     SFP+-10G-SR
      Xcvr 2      REV 01   740-021308   AD0915E003K      SFP+-10G-SR
      Xcvr 3      REV 01   740-021308   AD0915E003A      SFP+-10G-SR
    PIC 1                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80MRC          SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80NL5          SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80NKN          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80N3U          SFP+-10G-SR
    PIC 2                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80N1T          SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AJ808DJ          SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80NG4          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80FND          SFP+-10G-SR
    PIC 3                  BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80FKQ          SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80NLT          SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80NKR          SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80LKM          SFP+-10G-SR
```

```
FPC 10           REV 32   750-028467   ABBN6813         MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBK6542         AMPC PMB
  PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80NA3          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80NLF          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80MRH          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80KE4          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   973152A00030     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80L9H          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80ME8          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80NLR          SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80NG1          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80MCA          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80LFC          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80LEM          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80N9X          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80LAC          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80LF2          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80N8T          SFP+-10G-SR
FPC 11           REV 30   750-028467   ABBN0281         MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBN0526         AMPC PMB
  PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11F01326        SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E03973        SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E00950        SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E00674        SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E00775        SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E04461        SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E01074        SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E02821        SFP+-10G-USR
  PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E04501        SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E00757        SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01623        SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E01022        SFP+-10G-USR
  PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E04359        SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E02751        SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E02736        SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E01178        SFP+-10G-USR
FPC 12           REV 32   750-028467   ABBN6796         MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBN7259         AMPC PMB
  PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   B11K01856        SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11K01853        SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11K01863        SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02863     SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02668     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02881     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A01671     SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02627     SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02725     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02692     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02730     SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A03081     SFP+-10G-SR
```



```
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   163363A02736      SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02568      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A02747      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   163363A02579      SFP+-10G-SR
FPC 13          REV 30   750-028467   ABBN0270          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBJ0966          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80NL1           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   AK80NXW           SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   AK80KD2           SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   AK80FMD           SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80NKQ           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   AK80MGH           SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   AK80N38           SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   AK80NL7           SFP+-10G-SR
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80LEL           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   AK80NKD           SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   AK80KCY           SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   AK80LHK           SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80M5J           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   AK80MBE           SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   AK80NLG           SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   AK80LFH           SFP+-10G-SR
FPC 14          REV 32   750-028467   ABBN6790          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBK6515          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80LZM           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   AK80MCC           SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   AK80KCM           SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   AK80KE0           SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-021310   C10F99155         SFP+-10G-LRM
    Xcvr 1      REV 01   740-021310   C10F99049         SFP+-10G-LRM
    Xcvr 2      REV 01   740-021310   C10F99128         SFP+-10G-LRM
    Xcvr 3      REV 01   740-021310   C10F99169         SFP+-10G-LRM
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80LF3           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02597      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A03060      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   163363A03057      SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   AK80LEX           SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   AK80FEU           SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   AK80FNM           SFP+-10G-SR
    Xcvr 3      REV 01   740-021308   AJQ0Q5G           SFP+-10G-SR
FPC 15          REV 32   750-028467   ABBN6791          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBN7289          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   B11K00424         SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   B11K01849         SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   B11K01862         SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   B11K01852         SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   B11K00427         SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   B11K00430         SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   B11K01854         SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   B11K00426         SFP+-10G-SR
```

```
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   B11K00429         SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   B11K01864         SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   B11K01850         SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   B11K00522         SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E01144         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E00985         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E00796         SFP+-10G-USR
    Xcvr 3      REV 01   740-031980   B11K01866         SFP+-10G-SR
FPC 16          REV 30   750-028467   ABBM4592          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBN0465          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11F01435         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E01052         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11F01328         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11F01254         SFP+-10G-USR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E02738         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E02881         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11F01624         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E00889         SFP+-10G-USR
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E02883         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E00681         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E04306         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E02813         SFP+-10G-USR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11F01801         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E02753         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E01156         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E04324         SFP+-10G-USR
FPC 17          REV 32   750-028467   ABBN6810          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBN7237          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   163363A02638      SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02082      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A01674      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   163363A03058      SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   163363A03048      SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02729      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A02566      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   163363A02567      SFP+-10G-SR
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   163363A02878      SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02739      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A01959      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   163363A02660      SFP+-10G-SR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   163363A02731      SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02588      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A02673      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   163363A02654      SFP+-10G-SR
FPC 18          REV 30   750-028467   ABBM4739          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBN0487          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   163363A02569      SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   163363A02886      SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   163363A03082      SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   133363A00297      SFP+-10G-SR
```



```
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02726      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A03050      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02884      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A03076      SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02581      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02873      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02582      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A03083      SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031981   UL70BU6           SFP+-10G-LR
    Xcvr 1       REV 01   740-031981   UL50QC6           SFP+-10G-LR
    Xcvr 2       REV 01   740-031981   UL708N6           SFP+-10G-LR
    Xcvr 3       REV 01   740-031981   UL603KK           SFP+-10G-LR
FPC 19           REV 32   750-028467   ABBN6827          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBK6508          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A01688      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A01724      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A01773      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02593      SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A03061      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A03056      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02669      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A03070      SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02572      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02697      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02585      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A03052      SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02591      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02649      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02577      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02698      SFP+-10G-SR
ADC 0            REV 13   750-043596   ABBX5561          Adapter Card
ADC 1            REV 13   750-043596   ABBX5546          Adapter Card
ADC 2            REV 13   750-043596   ABBX5535          Adapter Card
ADC 3            REV 13   750-043596   ABBX5552          Adapter Card
ADC 4            REV 13   750-043596   ABBX5581          Adapter Card
ADC 5            REV 13   750-043596   ABBX5545          Adapter Card
ADC 6            REV 13   750-043596   ABBX5554          Adapter Card
ADC 7            REV 07   750-043596   ABBV7194          Adapter Card
ADC 8            REV 07   750-043596   ABBV7251          Adapter Card
ADC 9            REV 07   750-043596   ABBV7202          Adapter Card
ADC 10           REV 13   750-043596   ABBX5538          Adapter Card
ADC 11           REV 13   750-043596   ABBX5566          Adapter Card
ADC 12           REV 13   750-043596   ABBX5542          Adapter Card
ADC 13           REV 13   750-043596   ABBX5539          Adapter Card
ADC 14           REV 13   750-043596   ABBX5555          Adapter Card
ADC 15           REV 13   750-043596   ABBX5557          Adapter Card
ADC 16           REV 13   750-043596   ABBX5536          Adapter Card
ADC 17           REV 13   750-043596   ABBX5559          Adapter Card
ADC 18           REV 13   750-043596   ABBX5537          Adapter Card
ADC 19           REV 11   750-043596   ABBW5685          Adapter Card
Fan Tray 0       REV 2A   760-046960   ACAY0030          172mm FanTray - 6 Fans
Fan Tray 1       REV 2A   760-046960   ACAY0039          172mm FanTray - 6 Fans
Fan Tray 2       REV 2A   760-046960   ACAY0033          172mm FanTray - 6 Fans
Fan Tray 3       REV 2A   760-046960   ACAY0062          172mm FanTray - 6 Fans
```


show chassis hardware detail (MX2020 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11E2227AFJ      MX2020
Midplane         REV 27   750-040240   ABAB9384          Lower Power Midplane
Midplane 1       REV 04   711-032386   ABAB9386          Upper Backplane
PMP 1            REV 05   711-032428   ACAJ1821          Upper Power Midplane
PMP 0            REV 04   711-032426   ACAJ1524          Lower Power Midplane
FPM Board        REV 06   760-040242   ABBT8837          Front Panel Display
PSM 0            REV 01   740-045050   1E02224006G       DC 52V Power Supply
Module
PSM 1            REV 01   740-045050   1E022240053       DC 52V Power Supply
Module
PSM 2            REV 01   740-045050   1E02224004K       DC 52V Power Supply
Module
PSM 3            REV 01   740-045050   1E022240056       DC 52V Power Supply
Module
PSM 4            REV 01   740-045050   1E022240054       DC 52V Power Supply
Module
PSM 5            REV 01   740-045050   1E02224005H       DC 52V Power Supply
Module
PSM 6            REV 01   740-045050   1E02224006S       DC 52V Power Supply
Module
PSM 7            REV 01   740-045050   1E02224005M       DC 52V Power Supply
Module
PSM 8            REV 01   740-045050   1E022240062       DC 52V Power Supply
Module
PSM 9            REV 03   740-045050   1EDB2350095       DC 52V Power Supply
Module
PSM 10           REV 03   740-045050   1EDB235009L       DC 52V Power Supply
Module
PSM 11           REV 03   740-045050   1EDB2350092       DC 52V Power Supply
Module
PSM 12           REV 03   740-045050   1EDB23500AT       DC 52V Power Supply
Module
PSM 13           REV 03   740-045050   1EDB2350094       DC 52V Power Supply
Module
PSM 15           REV 03   740-045050   1EDB235008X       DC 52V Power Supply
Module
PDM 0            REV 01   740-045234   1E012150033       DC Power Dist Module
PDM 1            REV 01   740-045234   1E012150027       DC Power Dist Module
PDM 2            REV 01   740-045234   1E262250072       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009094138        RE-S-1800x4
  ad0    3998 MB  Virtium - TuffDisk VCF3 20110825A021D0000064 Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-499551-000347 Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009089709        RE-S-1800x4
  ad0    3831 MB  UGB30SFA4000T1       SFA4000T1 00000113 Compact Flash
  ad1   30533 MB  UGB94ARF32H0S3-KC    UNIGEN-478612-001044 Disk 1
CB 0             REV 08   750-040257   CAAB3482          Control Board
CB 1             REV 04   750-040257   ZT2864            Control Board
SPMB 0           REV 02   711-041855   CAAA6141          PMB Board
SPMB 1           REV 01   711-041855   ZS2275            PMB Board
```



```
SFB 0           REV 05   711-044466   ABBT2161          Switch Fabric Board
SFB 1           REV 05   711-044466   ABBT2159          Switch Fabric Board
SFB 2           REV 05   711-044466   ABBX3718          Switch Fabric Board
SFB 3           REV 05   711-044466   ABBT2152          Switch Fabric Board
SFB 4           REV 05   711-044466   ABBT2160          Switch Fabric Board
SFB 5           REV 05   711-044466   ABBT2145          Switch Fabric Board
SFB 6           REV 05   711-044466   ABBT2150          Switch Fabric Board
SFB 7           REV 05   711-044466   ABBT2163          Switch Fabric Board
FPC 0           REV 30   750-028467   ABBN0284          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBN0507          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E00990         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E04357         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11F01327         SFP+-10G-USR
Xcvr 3       REV 01   740-030658   B11E04375         SFP+-10G-USR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E02760         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E02904         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E03963         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E00756         SFP+-10G-USR
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E04418         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E01077         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E01128         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11F01253         SFP+-10G-USR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E01140         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11F01626         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E01075         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E01177         SFP+-10G-USR
FPC 1           REV 30   750-028467   ABBN0308          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBJ1095          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E04305         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E01147         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E01195         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11F01743         SFP+-10G-USR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11F01892         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E02880         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E00725         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E01057         SFP+-10G-USR
  PIC 2                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11E02816         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11C04501         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E02764         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E00789         SFP+-10G-USR
  PIC 3                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-030658   B11F01250         SFP+-10G-USR
    Xcvr 1      REV 01   740-030658   B11E02847         SFP+-10G-USR
    Xcvr 2      REV 01   740-030658   B11E00787         SFP+-10G-USR
    Xcvr 3      REV 01   740-030658   B11E03803         SFP+-10G-USR
FPC 2           REV 30   750-028467   ABBN0316          MPC 3D 16x 10GE
  CPU           REV 10   711-029089   ABBJ1082          AMPC PMB
  PIC 0                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   B11K00523         SFP+-10G-SR
    Xcvr 1      REV 01   740-031980   B11K01848         SFP+-10G-SR
    Xcvr 2      REV 01   740-031980   B11K01865         SFP+-10G-SR
    Xcvr 3      REV 01   740-031980   B11K00540         SFP+-10G-SR
  PIC 1                  BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0      REV 01   740-031980   B11K00422         SFP+-10G-SR
```



```
      Xcvr 1       REV 01   740-031980   B11K00428         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K00423         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K01855         SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   B11K01847         SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11K00526         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K00529         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K00525         SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
  Xcvr 0       REV 01   740-031980   B11K00425         SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11K00530         SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   B11K01851         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11K00528         SFP+-10G-SR
  FPC 3            REV 32   750-028467   ABBN6832          MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBK6534          AMPC PMB
    PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80MB4           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80FQ6           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80N1F           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80NLQ           SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80KDR           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80FGJ           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80N5G           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80KD8           SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80LET           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80N1X           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NRF           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80NL2           SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80N3D           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80MRB           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80LEQ           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80LER           SFP+-10G-SR
  FPC 4            REV 32   750-028467   ABBN6811          MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBN7288          AMPC PMB
    PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80NK8           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80LJG           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80LBU           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80N21           SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80LEU           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80NLM           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NL6           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80LES           SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80LEN           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80ME0           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80LMG           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80MM1           SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80MG7           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80KF9           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NRQ           SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80NLE           SFP+-10G-SR
  FPC 5            REV 32   750-028467   ABBN6791          MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBN7289          AMPC PMB
    PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   B11K00424         SFP+-10G-SR
```

```
    Xcvr 1       REV 01   740-031980   B11K01849         SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11K01862         SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11K01852         SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP
Xcvr 0       REV 01   740-031980   B11K00427         SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11K00430         SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11K01854         SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11K00426         SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   B11K00429         SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11K01864         SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11K01850         SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11K00522         SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E01144         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E00985         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E00796         SFP+-10G-USR
    Xcvr 3       REV 01   740-031980   B11K01866         SFP+-10G-SR
FPC 6            REV 30   750-028467   ABBM4592          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBN0465          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11F01435         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E01052         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01328         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11F01254         SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E02738         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E02881         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11F01624         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E00889         SFP+-10G-USR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11E02883         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E00681         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E04306         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E02813         SFP+-10G-USR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-030658   B11F01801         SFP+-10G-USR
    Xcvr 1       REV 01   740-030658   B11E02753         SFP+-10G-USR
    Xcvr 2       REV 01   740-030658   B11E01156         SFP+-10G-USR
    Xcvr 3       REV 01   740-030658   B11E04324         SFP+-10G-USR
FPC 7            REV 32   750-028467   ABBN6810          MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ABBN7237          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A03058      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02082      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A01674      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02638      SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A03048      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02729      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A02566      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02567      SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02878      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02739      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   163363A01959      SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   163363A02660      SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   163363A02731      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   163363A02588      SFP+-10G-SR
 Xcvr 2       REV 01   740-031980   163363A02673      SFP+-10G-SR
```

```
      Xcvr 3        REV 01   740-031980   163363A02654      SFP+-10G-SR
  FPC 8             REV 30   750-028467   ABBM4739          MPC 3D 16x 10GE
    CPU             REV 10   711-029089   ABBN0487          AMPC PMB
    PIC 0                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02569      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02886      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A03082      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   133363A00297      SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02726      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A03050      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02884      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A03076      SFP+-10G-SR
    PIC 2                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02581      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02873      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02582      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A03083      SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031981   UL70BU6           SFP+-10G-LR
      Xcvr 1        REV 01   740-031981   UL50QC6           SFP+-10G-LR
      Xcvr 2        REV 01   740-031981   UL708N6           SFP+-10G-LR
      Xcvr 3        REV 01   740-031981   UL603KK           SFP+-10G-LR
  FPC 9             REV 32   750-028467   ABBN6827          MPC 3D 16x 10GE
    CPU             REV 10   711-029089   ABBK6508          AMPC PMB
    PIC 0                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A01688      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A01724      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A01773      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A02593      SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A03061      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A03056      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02669      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A03070      SFP+-10G-SR
    PIC 2                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02572      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02697      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02585      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A03052      SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02591      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02649      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02577      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A02698      SFP+-10G-SR
  FPC 10            REV 30   750-028467   ABBN0302          MPC 3D 16x 10GE
    CPU             REV 10   711-029089   ABBN0495          AMPC PMB
    PIC 0                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11F01581         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11E01176         SFP+-10G-USR
      Xcvr 2        REV 01   740-030658   B11F01251         SFP+-10G-USR
      Xcvr 3        REV 01   740-030658   B11E02752         SFP+-10G-USR
    PIC 1                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11E00786         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11E01020         SFP+-10G-USR
   Xcvr 2        REV 01   740-030658   B11E01023         SFP+-10G-USR
      Xcvr 3        REV 01   740-030658   B11E02819         SFP+-10G-USR
    PIC 2                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11E02812         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11D04437         SFP+-10G-USR
      Xcvr 2        REV 01   740-030658   B11F01279         SFP+-10G-USR
```



```
      Xcvr 3      REV 01   740-030658   B11F01333      SFP+-10G-USR
    PIC 3                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-030658   B11E00978      SFP+-10G-USR
      Xcvr 1      REV 01   740-030658   B11E01018      SFP+-10G-USR
      Xcvr 2      REV 01   740-030658   B11F01784      SFP+-10G-USR
      Xcvr 3      REV 01   740-031980   AK80NKP        SFP+-10G-SR
  FPC 11          REV 32   750-028467   ABBN6790       MPC 3D 16x 10GE
    CPU           REV 10   711-029089   ABBK6515       AMPC PMB
    PIC 0                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80LZM        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80MCC        SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80KCM        SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80KE0        SFP+-10G-SR
    PIC 1                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-021310   C10F99155      SFP+-10G-LRM
      Xcvr 1      REV 01   740-021310   C10F99049      SFP+-10G-LRM
      Xcvr 2      REV 01   740-021310   C10F99128      SFP+-10G-LRM
      Xcvr 3      REV 01   740-021310   C10F99169      SFP+-10G-LRM
    PIC 2                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80LF3        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   163363A02597   SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   163363A03060   SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   163363A03057   SFP+-10G-SR
    PIC 3                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80LEX        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80FEU        SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80FNM        SFP+-10G-SR
      Xcvr 3      REV 01   740-021308   AJQ0Q5G        SFP+-10G-SR
  FPC 12          REV 30   750-028467   ZM5111         MPC 3D 16x 10GE
    CPU           REV 10   711-029089   ZP6607         AMPC PMB
    PIC 0                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80LJA        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80MFZ        SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80NKL        SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80KF4        SFP+-10G-SR
    PIC 1                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80FBJ        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80MM2        SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80LJV        SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80NXV        SFP+-10G-SR
    PIC 2                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80N1H        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80NLS        SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80FL5        SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80NL9        SFP+-10G-SR
    PIC 3                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-031980   AK80NG2        SFP+-10G-SR
      Xcvr 1      REV 01   740-031980   AK80KDU        SFP+-10G-SR
      Xcvr 2      REV 01   740-031980   AK80MG1        SFP+-10G-SR
      Xcvr 3      REV 01   740-031980   AK80MM0        SFP+-10G-SR
  FPC 13          REV 30   750-028467   ABBN0208       MPC 3D 16x 10GE
    CPU           REV 10   711-029089   ABBJ1084       AMPC PMB
    PIC 0                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-030658   B11E04745      SFP+-10G-USR
      Xcvr 1      REV 01   740-030658   B11F01570      SFP+-10G-USR
      Xcvr 2      REV 01   740-030658   B11E04388      SFP+-10G-USR
      Xcvr 3      REV 01   740-030658   B11F01439      SFP+-10G-USR
    PIC 1                  BUILTIN      BUILTIN        4x 10GE(LAN) SFP+
      Xcvr 0      REV 01   740-030658   B11E04739      SFP+-10G-USR
      Xcvr 1      REV 01   740-030658   B11F01869      SFP+-10G-USR
      Xcvr 2      REV 01   740-030658   B11F01675      SFP+-10G-USR
```



```
      Xcvr 3       REV 01   740-030658   B11F01901        SFP+-10G-USR
    PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-030658   B11F01346        SFP+-10G-USR
      Xcvr 1       REV 01   740-030658   B11F01288        SFP+-10G-USR
      Xcvr 2       REV 01   740-030658   B11F01824        SFP+-10G-USR
      Xcvr 3       REV 01   740-030658   B11E04312        SFP+-10G-USR
    PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-030658   B11E02811        SFP+-10G-USR
      Xcvr 1       REV 01   740-030658   B11E03847        SFP+-10G-USR
      Xcvr 2       REV 01   740-030658   B11F01495        SFP+-10G-USR
      Xcvr 3       REV 01   740-030658   B11F01265        SFP+-10G-USR
  FPC 14           REV 23   750-028467   YN2977           MPC 3D 16x 10GE
    CPU            REV 10   711-029089   YP1856           AMPC PMB
    PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   183363A00875     SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   183363A00851     SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   183363A00772     SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   183363A00882     SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   183363A00735     SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   183363A00169     SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   183363A00726     SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   183363A00077     SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   183363A00168     SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   183363A00676     SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   183363A00732     SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   183363A00091     SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   183363A00725     SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   183363A00642     SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   183363A00871     SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   183363A00853     SFP+-10G-SR
  FPC 15           REV 32   750-028467   ABBN6798         MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBK6556         AMPC PMB
    PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   9ZDZ06A00055     SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   183363A00239     SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   AD0915E003K      SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   AD0915E003A      SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80MRC          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80NL5          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NKN          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80N3U          SFP+-10G-SR
   PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80N1T          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AJ808DJ          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NG4          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80FND          SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80FKQ          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80NLT          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80NKR          SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   AK80LKM          SFP+-10G-SR
  FPC 16           REV 30   750-028467   ABBN0270         MPC 3D 16x 10GE
    CPU            REV 10   711-029089   ABBJ0966         AMPC PMB
    PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-031980   AK80NL1          SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   AK80NXW          SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AK80KD2          SFP+-10G-SR
```



```
      Xcvr 3        REV 01   740-031980   AK80FMD           SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   AK80NKQ           SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   AK80MGH           SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   AK80N38           SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   AK80NL7           SFP+-10G-SR
    PIC 2                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   AK80M5J           SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   AK80NKD           SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   AK80KCY           SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   AK80LHK           SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   AK80LEL           SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   AK80MBE           SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   AK80NLG           SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   AK80LFH           SFP+-10G-SR
  FPC 17            REV 32   750-028467   ABBN6796          MPC 3D 16x 10GE
    CPU             REV 10   711-029089   ABBN7259          AMPC PMB
    PIC 0                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   B11K01856         SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   B11K01853         SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   B11K01863         SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A02863      SFP+-10G-SR
    PIC 1                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02668      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02881      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A01671      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A02627      SFP+-10G-SR
    PIC 2                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02725      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02692      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02730      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A03081      SFP+-10G-SR
    PIC 3                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-031980   163363A02736      SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   163363A02568      SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   163363A02747      SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   163363A02579      SFP+-10G-SR
  FPC 18            REV 30   750-028467   ABBN0281          MPC 3D 16x 10GE
    CPU             REV 10   711-029089   ABBN0526          AMPC PMB
  PIC 0                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11F01326         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11E03973         SFP+-10G-USR
      Xcvr 2        REV 01   740-030658   B11E00950         SFP+-10G-USR
      Xcvr 3        REV 01   740-030658   B11E00674         SFP+-10G-USR
    PIC 1                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11E00775         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11E04461         SFP+-10G-USR
      Xcvr 2        REV 01   740-030658   B11E01074         SFP+-10G-USR
      Xcvr 3        REV 01   740-030658   B11E02821         SFP+-10G-USR
    PIC 2                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11E04501         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11E00757         SFP+-10G-USR
      Xcvr 2        REV 01   740-030658   B11F01623         SFP+-10G-USR
      Xcvr 3        REV 01   740-030658   B11E01022         SFP+-10G-USR
    PIC 3                    BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0        REV 01   740-030658   B11E04359         SFP+-10G-USR
      Xcvr 1        REV 01   740-030658   B11E02751         SFP+-10G-USR
      Xcvr 2        REV 01   740-030658   B11E02736         SFP+-10G-USR
      Xcvr 3        REV 01   740-030658   B11E01178         SFP+-10G-USR
  FPC 19            REV 32   750-028467   ABBN6813          MPC 3D 16x 10GE
```

```
  CPU            REV 10   711-029089   ABBK6542          AMPC PMB
  PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80NA3           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80NLF           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80MRH           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80KE4           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   973152A00030      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80L9H           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80ME8           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80NLR           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80NG1           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80MCA           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80LFC           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80LEM           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80N9X           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AK80LAC           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80LF2           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AK80N8T           SFP+-10G-SR
ADC 0            REV 13   750-043596   ABBX5561          Adapter Card
ADC 1            REV 13   750-043596   ABBX5546          Adapter Card
ADC 2            REV 13   750-043596   ABBX5535          Adapter Card
ADC 3            REV 13   750-043596   ABBX5552          Adapter Card
ADC 4            REV 13   750-043596   ABBX5581          Adapter Card
ADC 5            REV 13   750-043596   ABBX5545          Adapter Card
ADC 6            REV 13   750-043596   ABBX5554          Adapter Card
ADC 7            REV 07   750-043596   ABBV7194          Adapter Card
ADC 8            REV 07   750-043596   ABBV7251          Adapter Card
ADC 9            REV 07   750-043596   ABBV7202          Adapter Card
ADC 10           REV 13   750-043596   ABBX5579          Adapter Card
ADC 11           REV 13   750-043596   ABBX5548          Adapter Card
ADC 12           REV 13   750-043596   ABBX5575          Adapter Card
ADC 13           REV 13   750-043596   ABBX5539          Adapter Card
ADC 14           REV 13   750-043596   ABBX5555          Adapter Card
ADC 15           REV 13   750-043596   ABBX5557          Adapter Card
ADC 16           REV 13   750-043596   ABBX5536          Adapter Card
ADC 17           REV 13   750-043596   ABBX5559          Adapter Card
ADC 18           REV 13   750-043596   ABBX5537          Adapter Card
ADC 19           REV 11   750-043596   ABBW5685          Adapter Card
Fan Tray 0       REV 04   760-046960   ACAY0090          172mm FanTray - 6 Fans
Fan Tray 1       REV 04   760-046960   ACAY0088          172mm FanTray - 6 Fans
Fan Tray 2       REV 04   760-046960   ACAY0089          172mm FanTray - 6 Fans
Fan Tray 3       REV 04   760-046960   ACAY0108          172mm FanTray - 6 Fans
```

show chassis hardware models (MX2020 Router)

```
user@host > show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 27   750-040240   ABAB9384          750-040240
FPM Board        REV 06   760-040242   ABBT8837          760-040242
PSM 0            REV 01   740-045050   1E02224006G       MX2000-PSM-HC-DC-S-A
PSM 1            REV 01   740-045050   1E022240053       MX2000-PSM-HC-DC-S-A
PSM 2            REV 01   740-045050   1E02224004K       MX2000-PSM-HC-DC-S-A
PSM 3            REV 01   740-045050   1E022240056       MX2000-PSM-HC-DC-S-A
PSM 4            REV 01   740-045050   1E022240054       MX2000-PSM-HC-DC-S-A
PSM 5            REV 01   740-045050   1E02224005H       MX2000-PSM-HC-DC-S-A
PSM 6            REV 01   740-045050   1E02224006S       MX2000-PSM-HC-DC-S-A
PSM 7            REV 01   740-045050   1E02224005M       MX2000-PSM-HC-DC-S-A
```



```
PSM 8            REV 01   740-045050   1E022240062       MX2000-PSM-HC-DC-S-A
PSM 9            REV 03   740-045050   1EDB2350095       MX2000-PSM-DC-S-A
PSM 10           REV 03   740-045050   1EDB235009L       MX2000-PSM-DC-S-A
PSM 11           REV 03   740-045050   1EDB2350092       MX2000-PSM-DC-S-A
PSM 12           REV 03   740-045050   1EDB23500AT       MX2000-PSM-DC-S-A
PSM 13           REV 03   740-045050   1EDB2350094       MX2000-PSM-DC-S-A
PSM 15           REV 03   740-045050   1EDB235008X       MX2000-PSM-DC-S-A
PDM 0            REV 01   740-045234   1E012150033
PDM 1            REV 01   740-045234   1E012150027
PDM 2            REV 01   740-045234   1E262250072       MX2000-PDM-DC-S-A
Routing Engine 0 REV 02   740-041821   9009094138        RE-S-1800X4-16G-S
Routing Engine 1 REV 02   740-041821   9009089709        RE-S-1800X4-16G-S
CB 0             REV 08   750-040257   CAAB3482          750-040257
CB 1             REV 04   750-040257   ZT2864            750-040257
SFB 0            REV 05   711-044466   ABBT2161          MX2000-SFB-S
SFB 1            REV 05   711-044466   ABBT2159          MX2000-SFB-S
SFB 2            REV 05   711-044466   ABBX3718          MX2000-SFB-S
SFB 4            REV 05   711-044466   ABBT2160          MX2000-SFB-S
SFB 5            REV 05   711-044466   ABBT2145          MX2000-SFB-S
SFB 7            REV 05   711-044466   ABBT2163          MX2000-SFB-S
FPC 0            REV 30   750-028467   ABBN0284          MPC-3D-16XGE-SFPP
FPC 1            REV 30   750-028467   ABBN0308          MPC-3D-16XGE-SFPP
FPC 2            REV 30   750-028467   ABBN0316          MPC-3D-16XGE-SFPP
FPC 3            REV 32   750-028467   ABBN6832          MPC-3D-16XGE-SFPP
FPC 4            REV 32   750-028467   ABBN6811          MPC-3D-16XGE-SFPP
FPC 5            REV 32   750-028467   ABBN6791          MPC-3D-16XGE-SFPP
FPC 6            REV 30   750-028467   ABBM4592          MPC-3D-16XGE-SFPP
FPC 7            REV 32   750-028467   ABBN6810          MPC-3D-16XGE-SFPP
FPC 8            REV 30   750-028467   ABBM4739          MPC-3D-16XGE-SFPP
FPC 9            REV 32   750-028467   ABBN6827          MPC-3D-16XGE-SFPP
FPC 10           REV 30   750-028467   ABBN0302          MPC-3D-16XGE-SFPP
FPC 11           REV 32   750-028467   ABBN6790          MPC-3D-16XGE-SFPP
FPC 12           REV 30   750-028467   ZM5111            MPC-3D-16XGE-SFPP
FPC 13           REV 30   750-028467   ABBN0208          MPC-3D-16XGE-SFPP
FPC 14           REV 23   750-028467   YN2977            MPC-3D-16XGE-SFPP
FPC 15           REV 32   750-028467   ABBN6798          MPC-3D-16XGE-SFPP
FPC 16           REV 30   750-028467   ABBN0270          MPC-3D-16XGE-SFPP
FPC 17           REV 32   750-028467   ABBN6796          MPC-3D-16XGE-SFPP
FPC 18           REV 30   750-028467   ABBN0281          MPC-3D-16XGE-SFPP
FPC 19           REV 32   750-028467   ABBN6813          MPC-3D-16XGE-SFPP
ADC 0            REV 13   750-043596   ABBX5561          PROTO-ASSEMBLY
ADC 1            REV 13   750-043596   ABBX5546          PROTO-ASSEMBLY
ADC 2            REV 13   750-043596   ABBX5535          MX2000-LC-ADAPTER
ADC 3            REV 13   750-043596   ABBX5552          MX2000-LC-ADAPTER
ADC 4            REV 13   750-043596   ABBX5581          MX2000-LC-ADAPTER
ADC 5            REV 13   750-043596   ABBX5545          PROTO-ASSEMBLY
ADC 6            REV 13   750-043596   ABBX5554          PROTO-ASSEMBLY
ADC 7            REV 07   750-043596   ABBV7194          MX2000-LC-ADAPTER
ADC 8            REV 07   750-043596   ABBV7251          MX2000-LC-ADAPTER
ADC 9            REV 07   750-043596   ABBV7202          MX2000-LC-ADAPTER
ADC 10           REV 13   750-043596   ABBX5579          MX2000-LC-ADAPTER
ADC 12           REV 13   750-043596   ABBX5575          MX2000-LC-ADAPTER
ADC 13           REV 13   750-043596   ABBX5539          PROTO-ASSEMBLY
ADC 14           REV 13   750-043596   ABBX5555          PROTO-ASSEMBLY
ADC 15           REV 13   750-043596   ABBX5557          MX2000-LC-ADAPTER
ADC 16           REV 13   750-043596   ABBX5536          PROTO-ASSEMBLY
ADC 17           REV 13   750-043596   ABBX5559          PROTO-ASSEMBLY
ADC 18           REV 13   750-043596   ABBX5537          PROTO-ASSEMBLY
ADC 19           REV 11   750-043596   ABBW5685          PROTO-ASSEMBLY
Fan Tray 0       REV 04   760-046960   ACAY0090
Fan Tray 1       REV 04   760-046960   ACAY0088
```

```
Fan Tray 2        REV 04    760-046960    ACAY0089
Fan Tray 3        REV 04    760-046960    ACAY0108
```

show chassis hardware clei-models (MX2020 Router)

```
user@ host > show chassis hardware clei-models
Hardware inventory:
Item              Version   Part number   CLEI code         FRU model number
Midplane          REV 27    750-040240    PROTOXCLEI        750-040240
FPM Board         REV 06    760-040242    PROTOXCLEI        760-040242
PSM 0             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 1             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 2             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 3             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 4             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 5             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 6             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 7             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 8             REV 01    740-045050    IPUPAJMKAA        MX2000-PSM-HC-DC-S-A
PSM 9             REV 03    740-045050    IPUPAJMKAA        MX2000-PSM-DC-S-A
PSM 10            REV 03    740-045050    IPUPAJMKAA        MX2000-PSM-DC-S-A
PSM 11            REV 03    740-045050    IPUPAJMKAA        MX2000-PSM-DC-S-A
PSM 12            REV 03    740-045050    IPUPAJMKAA        MX2000-PSM-DC-S-A
PSM 13            REV 03    740-045050    IPUPAJMKAA        MX2000-PSM-DC-S-A
PSM 15            REV 03    740-045050    IPUPAJMKAA        MX2000-PSM-DC-S-A
PDM 0             REV 01    740-045234
PDM 1             REV 01    740-045234
PDM 2             REV 01    740-045234    IPUPAJSKAA        MX2000-PDM-DC-S-A
Routing Engine 0  REV 02    740-041821                      RE-S-1800X4-16G-S
Routing Engine 1  REV 02    740-041821                      RE-S-1800X4-16G-S
CB 0              REV 08    750-040257    PROTOXCLEI        750-040257
CB 1              REV 04    750-040257    PROTOXCLEI        750-040257
SFB 0             REV 05    711-044466    IPUCBA6CAA        MX2000-SFB-S
SFB 1             REV 05    711-044466    IPUCBA6CAA        MX2000-SFB-S
SFB 2             REV 05    711-044466    IPUCBA6CAA        MX2000-SFB-S
SFB 4             REV 05    711-044466    IPUCBA6CAA        MX2000-SFB-S
SFB 5             REV 05    711-044466    IPUCBA6CAA        MX2000-SFB-S
SFB 7             REV 05    711-044466    IPUCBA6CAA        MX2000-SFB-S
FPC 0             REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 1             REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 2             REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 3             REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 4             REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 5             REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 6             REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 7             REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 8             REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 9             REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 10            REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 11            REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 12            REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 13            REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 14            REV 23    750-028467                      MPC-3D-16XGE-SFPP
FPC 15            REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 16            REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 17            REV 32    750-028467                      MPC-3D-16XGE-SFPP
FPC 18            REV 30    750-028467                      MPC-3D-16XGE-SFPP
FPC 19            REV 32    750-028467                      MPC-3D-16XGE-SFPP
ADC 0             REV 13    750-043596    PROTOXCLEI        PROTO-ASSEMBLY
ADC 1             REV 13    750-043596    PROTOXCLEI        PROTO-ASSEMBLY
ADC 2             REV 13    750-043596    IPUCBA8CAA        MX2000-LC-ADAPTER
```

```
ADC 3           REV 13   750-043596   IPUCBA8CAA       MX2000-LC-ADAPTER
ADC 4           REV 13   750-043596   IPUCBA8CAA       MX2000-LC-ADAPTER
ADC 5           REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 6           REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 7           REV 07   750-043596   PROTOXCLEI       MX2000-LC-ADAPTER
ADC 8           REV 07   750-043596   PROTOXCLEI       MX2000-LC-ADAPTER
ADC 9           REV 07   750-043596   PROTOXCLEI       MX2000-LC-ADAPTER
ADC 10          REV 13   750-043596   IPUCBA8CAA       MX2000-LC-ADAPTER
ADC 12          REV 13   750-043596   IPUCBA8CAA       MX2000-LC-ADAPTER
ADC 13          REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 14          REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 15          REV 13   750-043596   IPUCBA8CAA       MX2000-LC-ADAPTER
ADC 16          REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 17          REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 18          REV 13   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
ADC 19          REV 11   750-043596   PROTOXCLEI       PROTO-ASSEMBLY
Fan Tray 0      REV 04   760-046960
Fan Tray 1      REV 04   760-046960
Fan Tray 2      REV 04   760-046960
Fan Tray 3      REV 04   760-046960
```

show chassis hardware (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis hardware
Hardware inventory:
Item            Version  Part number  Serial number    Description
Chassis                               JN120BADBAFJ     MX2020
Midplane        REV 51   750-040240   ABAB9243         Lower Backplane
Midplane 1      REV 04   711-032386   ABAB9399         Upper Backplane
PMP 1           REV 05   711-032428   ACAJ2541         Upper Power Midplane
PMP 0           REV 04   711-032426   ACAJ2194         Lower Power Midplane
FPM Board       REV 13   760-040242   ABCA8835         Front Panel Display
PSM 0           REV 01   740-050037   1EDB32403L5      DC 52V Power Supply
Module
PSM 1           REV 01   740-050037   1EDB32403L3      DC 52V Power Supply
Module
PSM 2           REV 01   740-050037   1EDB32403KM      DC 52V Power Supply
Module
PSM 3           REV 01   740-050037   1EDB3130079      DC 52V Power Supply
Module
PSM 4           REV 01   740-050037   1EDB3130077      DC 52V Power Supply
Module
PSM 5           REV 01   740-050037   1EDB3130020      DC 52V Power Supply
Module
PSM 6           REV 01   740-050037   1EDB313009S      DC 52V Power Supply
Module
PSM 7           REV 01   740-050037   1EDB313008E      DC 52V Power Supply
Module
PSM 8           REV 01   740-050037   1EDB3130063      DC 52V Power Supply
Module
PSM 12          REV 01   740-050037   1EDB3130026      DC 52V Power Supply
Module
PSM 13          REV 01   740-050037   1EDB3130074      DC 52V Power Supply
Module
PSM 14          REV 01   740-050037   1EDB313009D      DC 52V Power Supply
Module
PSM 15          REV 01   740-050037   1EDB3130024      DC 52V Power Supply
Module
PSM 16          REV 01   740-050037   1EDB3130054      DC 52V Power Supply
Module
PSM 17          REV 01   740-050037   1EDB3130080      DC 52V Power Supply
```

```
Module
PDM 0            REV 03   740-045234   1EGA3170144       DC Power Dist Module
PDM 1            REV 03   740-045234   1EGA3170158       DC Power Dist Module
PDM 2            REV 03   740-045234   1EGA3170182       DC Power Dist Module
PDM 3            REV 03   740-045234   1EGA3170207       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009112112        RE-S-1800x4
Routing Engine 1 REV 02   740-041821   9009112087        RE-S-1800x4
CB 0             REV 23   750-040257   CABA2295          Control Board
CB 1             REV 23   750-040257   CABE8379          Control Board
SPMB 0           REV 02   711-041855   ABCE8851          PMB Board
SPMB 1           REV 02   711-041855   ABCE8839          PMB Board
SFB 0            REV 06   711-044466   ABCD5001          Switch Fabric Board
SFB 1            REV 06   711-044466   ABCD5034          Switch Fabric Board
SFB 2            REV 06   711-044466   ABCH3899          Switch Fabric Board
SFB 3            REV 06   711-044466   ABCD5020          Switch Fabric Board
SFB 4            REV 06   711-044466   ABCD4975          Switch Fabric Board
SFB 5            REV 06   711-044466   ABCH3881          Switch Fabric Board
SFB 6            REV 06   711-044466   ABCD5026          Switch Fabric Board
SFB 7            REV 06   711-044466   ABCD5032          Switch Fabric Board
FPC 0            REV 39   750-045715   CACD1902          MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 09   711-045719   CACB1933          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           12X10GE SFPP OTN
    Xcvr 0       REV 01   740-031980   B11F00361         SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   19T511101854      SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   19T511100377      SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   ANT0878           SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   19T511100398      SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AQ4363J           SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   19T511101377      SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   ANT072M           SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   AG90C7N           SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   AM30M09           SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   B10E01016         SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           12X10GE SFPP OTN
 Xcvr 0       REV 01   740-031980   B10L04151         SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   19T511101379      SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQ5036J           SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AG90C4M           SFP+-10G-SR
    Xcvr 4       REV 01   740-021308   19T511101104      SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   AQ502ZM           SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AN10KY2           SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   AQ43G41           SFP+-10G-SR
    Xcvr 8       REV 01   740-021308   AQ41F04           SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   AMS16N3           SFP+-10G-SR
    Xcvr 10      REV 01   740-021308   AMH04Y3           SFP+-10G-SR
    Xcvr 11      REV 01   740-021308   ANA093E           SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN           3X40GE QSFPP
  PIC 3                   BUILTIN      BUILTIN           3X40GE QSFPP
  WAN MEZZ       REV 09   750-049136   CABN0410          MPC5E 24XGE OTN Mezz
FPC 1            REV 11   750-045372   CABK8112          MPCE Type 3 3D
  CPU            REV 08   711-035209   CABJ6621          HMPC PMB 2G
  MIC 0          REV 07   750-033307   CAAZ2897          10X10GE SFPP
    PIC 0                 BUILTIN      BUILTIN           10X10GE SFPP
      Xcvr 0     REV 01   740-021308   AQ501VK           SFP+-10G-SR
      Xcvr 1     REV 01   740-021308   AQ501YC           SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   AQ43HJF           SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   AQ43H8D           SFP+-10G-SR
      Xcvr 4     REV 01   740-021308   19T511100370      SFP+-10G-SR
      Xcvr 5     REV 01   740-031980   153363A00763      SFP+-10G-SR
      Xcvr 6     REV 01   740-021308   APH2LXB           SFP+-10G-SR
      Xcvr 7     REV 01   740-031980   AMC0LVV           SFP+-10G-SR
```



```
        Xcvr 8     REV 01   740-031980   B11F00230         SFP+-10G-SR
    MIC 1          REV 14   750-033196   CAAP1390          1X100GE CXP
      PIC 2                 BUILTIN      BUILTIN           1X100GE CXP
        Xcvr 0     REV 01   740-032166   XB11F000M         CFP2-100G-SR10
  FPC 2            REV 17   750-037355   CAAS5826          MPC4E 3D 2CGE+8XGE
    CPU            REV 08   711-035209   CAAR3986          HMPC PMB 2G
    PIC 0                   BUILTIN      BUILTIN           4x10GE SFPP
      Xcvr 0       REV 01   740-021308   T09F43722         SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   ALP0KXF           SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   AQ502FG           SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   AQ502T7           SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           1X100GE CFP
      Xcvr 0       REV 01   740-035329   X12J00571         CFP-100G-SR10
    PIC 2                   BUILTIN      BUILTIN           4x10GE SFPP
      Xcvr 0       REV 01   740-031980   AJ71KEH           SFP+-10G-SR
      Xcvr 1       REV 01   740-031980   B11E01355         SFP+-10G-SR
      Xcvr 3       REV 01   740-031980   B11F00249         SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           1X100GE CFP
  FPC 3            REV 05   750-044444   CAAY9920          MPCE Type 2 3D P
    CPU            REV 04   711-038484   CAAW3639          MPCE PMB 2G
    MIC 0          REV 28   750-028387   CAAX1083          3D 4x 10GE  XFP
      PIC 0                 BUILTIN      BUILTIN           2x 10GE  XFP
        Xcvr 0              NON-JNPR     CC07BK05B         XFP-10G-SR
        Xcvr 1     REV 01   740-011571   C728XJ00U         XFP-10G-SR
      PIC 1                 BUILTIN      BUILTIN           2x 10GE  XFP
        Xcvr 0              NON-JNPR     T12L92339         XFP-10G-SR
    QXM 0          REV 06   711-028408   CAAW4915          MPC QXM
    QXM 1          REV 06   711-028408   CAAW4894          MPC QXM
  FPC 4            REV 18   750-046005   CACH5661          MPC5E 3D Q 2CGE+4XGE
   CPU           REV 09   711-045719   CACF2880          RMPC PMB
    PIC 0                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    PIC 1                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
      Xcvr 0       REV 01   740-046563   XD16FC03Y         CFP2-100G-SR10
    PIC 2                   BUILTIN      BUILTIN           2X10GE SFPP OTN
    PIC 3                   BUILTIN      BUILTIN           1X100GE CFP2 OTN
      Xcvr 0       REV 01   740-049775   J13K72997         CFP2-100G-LR4-D
  FPC 5            REV 35   750-028467   CAAR2623          MPC 3D 16x 10GE
    CPU            REV 11   711-029089   CAAR0491          AMPC PMB
    PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   AQ5027T           SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   AQ502J0           SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   AQ5027S           SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   AQ501Y7           SFP+-10G-SR
    PIC 1                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   AQ501YB           SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   AQ503EB           SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   AQ43HJH           SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   AQ43J0Y           SFP+-10G-SR
    PIC 2                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   AQ50352           SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   AQ501X6           SFP+-10G-SR
      Xcvr 2       REV 01   740-021308   AQ502NV           SFP+-10G-SR
      Xcvr 3       REV 01   740-021308   AQ502ZJ           SFP+-10G-SR
    PIC 3                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
      Xcvr 0       REV 01   740-021308   AQ502H4           SFP+-10G-SR
      Xcvr 1       REV 01   740-021308   AQ43HJK           SFP+-10G-SR
      Xcvr 2       REV 01   740-031980   AJ30CU7           SFP+-10G-SR
  FPC 9            REV 30   750-044130   ABCF5773          MPC6E 3D
    CPU            REV 09   711-045719   ABCF1270          RMPC PMB
    MIC 0          REV 05   750-049457   ABCD7829          2X100GE CFP2 OTN
      PIC 0                 BUILTIN      BUILTIN           2X100GE CFP2 OTN
```

```
      Xcvr 0              NON-JNPR     FE13F000K          CFP2-100G-SR10
      Xcvr 1     REV 01   740-048813   XD32FE017          CFP2-100G-LR-D
  MIC 1          REV 07   750-049457   ABCK2812           2X100GE CFP2 OTN
    PIC 1                 BUILTIN      BUILTIN            2X100GE CFP2 OTN
      Xcvr 0     REV 01   740-048813   XD32FE018          CFP2-100G-SR10
      Xcvr 1              NON-JNPR     FE13F000E          CFP2-100G-LR4-D
  XLM 0          REV 05.2.00 711-046638 ABCF5915          MPC6E XL
  XLM 1          REV 05.2.00 711-046638 ABCF5916          MPC6E XL
FPC 10           REV 36   750-044130   ABCS8602           MPC6E 3D
  CPU            REV 09   711-045719   ABCS8779           RMPC PMB
  MIC 0          REV 06   750-049979   ABCK2656           24X10GE SFPP OTN
    PIC 0                 BUILTIN      BUILTIN            24X10GE SFPP OTN
      Xcvr 0     REV 01   740-021308   AQ43J08            SFP+-10G-SR
      Xcvr 1     REV 01   740-021308   AQE1Y2E            SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   AQE1UW4            SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   AQE1MQF            SFP+-10G-SR
      Xcvr 4     REV 01   740-021308   AQG0MN1            SFP+-10G-SR
      Xcvr 5     REV 01   740-021308   AQE1L9M            SFP+-10G-SR
      Xcvr 6     REV 01   740-021308   AQG0MPD            SFP+-10G-SR
      Xcvr 7     REV 01   740-021308   AQE1Y2B            SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AQG0LT5            SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AQD2ET4            SFP+-10G-SR
      Xcvr 10    REV 01   740-021308   AQG0MPC            SFP+-10G-SR
      Xcvr 11    REV 01   740-021308   AQG0M63            SFP+-10G-SR
     Xcvr 12    REV 01   740-021308   AQG0LT1            SFP+-10G-SR
      Xcvr 13    REV 01   740-021308   AQG0M4L            SFP+-10G-SR
      Xcvr 14    REV 01   740-021308   AQG0LS7            SFP+-10G-SR
      Xcvr 15    REV 01   740-021308   AQE1MQB            SFP+-10G-SR
      Xcvr 16    REV 01   740-021308   AQG0LZP            SFP+-10G-SR
      Xcvr 17    REV 01   740-021308   AQE1LU9            SFP+-10G-SR
      Xcvr 18    REV 01   740-021308   AQG0MRZ            SFP+-10G-SR
      Xcvr 19    REV 01   740-021308   AQE1MQ9            SFP+-10G-SR
      Xcvr 20    REV 01   740-021308   AQG0LRX            SFP+-10G-SR
      Xcvr 21    REV 01   740-021308   AQE1UWD            SFP+-10G-SR
      Xcvr 22    REV 01   740-021308   AQG0LT4            SFP+-10G-SR
      Xcvr 23    REV 01   740-021308   AQE1MQL            SFP+-10G-SR
  MIC 1          REV 12   750-050008   ABCK5372           4X100GE CXP
    PIC 1                 BUILTIN      BUILTIN            4X100GE CXP
      Xcvr 3     REV 01   740-046563   XD16FC02Z          CFP2-100G-SR10
  XLM 0          REV 07.2.00 711-046638 ABCK3481          MPC6E XL
  XLM 1          REV 07.2.00 711-046638 ABCK4725          MPC6E XL
FPC 17           REV 28   750-044130   ABBZ3873           MPC6E 3D
  CPU            REV 08   711-045719   ABBZ3770           RMPC PMB
  MIC 0          REV 11   750-046535   ABCC7731           24X10GE SFPP
    PIC 0                 BUILTIN      BUILTIN            24X10GE SFPP
      Xcvr 1     REV 01   740-021308   APK0543            SFP+-10G-SR
      Xcvr 2     REV 01   740-031980   B10G01119          SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   AQ502SX            SFP+-10G-SR
      Xcvr 4     REV 01   740-021308   AQ43H84            SFP+-10G-SR
      Xcvr 5     REV 01   740-021308   AQ501TB            SFP+-10G-SR
      Xcvr 6     REV 01   740-021308   AQ502JZ            SFP+-10G-SR
      Xcvr 7     REV 01   740-021308   AQ502SC            SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AQ502JW            SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AQ502RM            SFP+-10G-SR
      Xcvr 10    REV 01   740-031980   AHK013B            SFP+-10G-SR
      Xcvr 11    REV 01   740-021308   AQG0MRT            SFP+-10G-SR
      Xcvr 13    REV 01   740-031980   AMC0JTC            SFP+-10G-SR
      Xcvr 14    REV 01   740-021308   ANA0MQ0            SFP+-10G-SR
      Xcvr 15    REV 01   740-021308   AQ502GS            SFP+-10G-SR
      Xcvr 16    REV 01   740-021308   AQG0M0J            SFP+-10G-SR
      Xcvr 17    REV 01   740-021308   AQG0MUR            SFP+-10G-SR
```



```
      Xcvr 18    REV 01   740-021308   AQG0MRR          SFP+-10G-SR
      Xcvr 19    REV 01   740-021308   AQG0M0F          SFP+-10G-SR
      Xcvr 20    REV 01   740-021308   AQ50312          SFP+-10G-SR
      Xcvr 21    REV 01   740-021308   AQ5032U          SFP+-10G-SR
      Xcvr 22    REV 01   740-021308   APE17B5          SFP+-10G-SR
      Xcvr 23    REV 01   740-021309   91D104A00011     SFP+-10G-LR
  MIC 1          REV 03   750-050008   ABCC4522         4X100GE CXP
    PIC 1                 BUILTIN      BUILTIN          4X100GE CXP
      Xcvr 0     REV 01   740-046563   XD16FC02U        CFP2-100G-SR10
      Xcvr 1     REV 01   740-046563   XC42FC03K        CFP2-100G-SR10
      Xcvr 2     REV 01   740-046563   XC42FC01Z        CFP2-100G-SR10
      Xcvr 3     REV 01   740-046563   XC42FC02U        CFP2-100G-SR10
  XLM 0          REV 04.2.00 711-046638 ABBZ3779        MPC6E XL
  XLM 1          REV 04.2.00 711-046638 ABBZ3780        MPC6E XL
FPC 18           REV 39   750-045715   CACD1910         MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 09   711-045719   CACD1817         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          12X10GE SFPP OTN
  PIC 1                   BUILTIN      BUILTIN          12X10GE SFPP OTN
  PIC 2                   BUILTIN      BUILTIN          3X40GE QSFPP
 Xcvr 0       REV 01   740-046565   QD130194         QSFP+-40G-SR4
    Xcvr 1       REV 01   740-046565   QD130193         QSFP+-40G-SR4
    Xcvr 2       REV 01   740-046565   QD130196         QSFP+-40G-SR4
  PIC 3                   BUILTIN      BUILTIN          3X40GE QSFPP
    Xcvr 0       REV 01   740-046565   QD130191         QSFP+-40G-SR4
    Xcvr 1       REV 01   740-046565   QD130198         QSFP+-40G-SR4
    Xcvr 2       REV 01   740-046565   QD130192         QSFP+-40G-SR4
  WAN MEZZ       REV 09   750-049136   CABN0411         MPC5E 24XGE OTN Mezz
FPC 19           REV 39   750-045715   CACD1908         MPC5E 3D Q 24XGE+6XLGE
  CPU            REV 09   711-045719   CACD1820         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          12X10GE SFPP OTN
    Xcvr 0       REV 01   740-021308   AQA0EXJ          SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQG0M6D          SFP+-10G-SR
    Xcvr 2       REV 01   740-021308   AQG0LW7          SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AQA0JKB          SFP+-10G-SR
    Xcvr 4       REV 01   740-021308   AQG0MTM          SFP+-10G-SR
    Xcvr 5       REV 01   740-021308   AQA07NE          SFP+-10G-SR
    Xcvr 6       REV 01   740-021308   AQG0M41          SFP+-10G-SR
    Xcvr 7       REV 01   740-021308   AQG0MU7          SFP+-10G-SR
    Xcvr 8       REV 01   740-021308   AQG0MUG          SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   AQG0MMX          SFP+-10G-SR
    Xcvr 10      REV 01   740-021308   AQG0M5K          SFP+-10G-SR
    Xcvr 11      REV 01   740-021308   AQG0LVZ          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          12X10GE SFPP OTN
  PIC 2                   BUILTIN      BUILTIN          3X40GE QSFPP
  PIC 3                   BUILTIN      BUILTIN          3X40GE QSFPP
    Xcvr 0       REV 01   740-046565   QD130242         QSFP+-40G-SR4
    Xcvr 1       REV 01   740-046565   QD130245         QSFP+-40G-SR4
    Xcvr 2       REV 01   740-046565   QD130613         QSFP+-40G-SR4
  WAN MEZZ       REV 09   750-049136   CABN0418         MPC5E 24XGE OTN Mezz
ADC 0            REV 17   750-043596   ABCD5378         Adapter Card
ADC 1            REV 17   750-043596   ABCD5465         Adapter Card
ADC 2            REV 17   750-043596   ABCD5431         Adapter Card
ADC 3            REV 17   750-043596   ABCD5356         Adapter Card
ADC 4            REV 02   750-043596   ZW1545           Adapter Card
ADC 5            REV 17   750-043596   ABCD5517         Adapter Card
ADC 18           REV 17   750-043596   ABCD5535         Adapter Card
ADC 19           REV 01   750-043596   ZV4127           Adapter Card
Fan Tray 0       REV 06   760-046960   ACAY0791         172mm FanTray - 6 Fans
Fan Tray 1       REV 06   760-046960   ACAY0788         172mm FanTray - 6 Fans
Fan Tray 2       REV 06   760-046960   ACAY0755         172mm FanTray - 6 Fans
Fan Tray 3       REV 06   760-046960   ACAY0441         172mm FanTray - 6 Fans
```


show chassis hardware detail (MX2020 Router with MPC5EQ and MPC6E)

```
user@host>show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN120BADBAFJ      MX2020
Midplane         REV 51   750-040240   ABAB9243          Lower Backplane
Midplane 1       REV 04   711-032386   ABAB9399          Upper Backplane
PMP 1            REV 05   711-032428   ACAJ2541          Upper Power Midplane
PMP 0            REV 04   711-032426   ACAJ2194          Lower Power Midplane
FPM Board        REV 13   760-040242   ABCA8835          Front Panel Display
PSM 0            REV 01   740-050037   1EDB32403L5       DC 52V Power Supply
Module
PSM 1            REV 01   740-050037   1EDB32403L3       DC 52V Power Supply
Module
PSM 2            REV 01   740-050037   1EDB32403KM       DC 52V Power Supply
Module
PSM 3            REV 01   740-050037   1EDB3130079       DC 52V Power Supply
Module
PSM 4            REV 01   740-050037   1EDB3130077       DC 52V Power Supply
Module
PSM 5            REV 01   740-050037   1EDB3130020       DC 52V Power Supply
Module
PSM 6            REV 01   740-050037   1EDB313009S       DC 52V Power Supply
Module
PSM 7            REV 01   740-050037   1EDB313008E       DC 52V Power Supply
Module
PSM 8            REV 01   740-050037   1EDB3130063       DC 52V Power Supply
Module
PSM 12           REV 01   740-050037   1EDB3130026       DC 52V Power Supply
Module
PSM 13           REV 01   740-050037   1EDB3130074       DC 52V Power Supply
Module
PSM 14           REV 01   740-050037   1EDB313009D       DC 52V Power Supply
Module
PSM 15           REV 01   740-050037   1EDB3130024       DC 52V Power Supply
Module
PSM 16           REV 01   740-050037   1EDB3130054       DC 52V Power Supply
Module
PSM 17           REV 01   740-050037   1EDB3130080       DC 52V Power Supply
Module
PDM 0            REV 03   740-045234   1EGA3170144       DC Power Dist Module
PDM 1            REV 03   740-045234   1EGA3170158       DC Power Dist Module
PDM 2            REV 03   740-045234   1EGA3170182       DC Power Dist Module
PDM 3            REV 03   740-045234   1EGA3170207       DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009112112        RE-S-1800x4
  ad0     3998 MB  Virtium - TuffDrive VCF P1T0200274310822 113 Compact Flash
  ad1    30533 MB  UGB94BPH32H0S1-KCI   11000031656       Disk 1
  usb0 (addr 1)  EHCI root hub 0       Intel             uhub0
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub1
  DIMM 0         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3         SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009112087        RE-S-1800x4
  ad0     3998 MB  Virtium - TuffDrive VCF P1T0200274310822 366 Compact Flash
  ad1    30533 MB  UGB94BPH32H0S1-KCI   11000039979       Disk 1
CB 0             REV 23   750-040257   CABA2295          Control Board
CB 1             REV 23   750-040257   CABE8379          Control Board
SPMB 0
```

```
SPMB 1
FPC 0            REV 39   750-045715   CACD1902         MPC5E 3D Q 24XGE+6XLGE
  CPU
FPC 1            REV 11   750-045372   CABK8112         MPCE Type 3 3D
  CPU
FPC 2            REV 17   750-037355   CAAS5826         MPC4E 3D 2CGE+8XGE
  CPU
FPC 3            REV 05   750-044444   CAAY9920         MPCE Type 2 3D P
  CPU
FPC 4            REV 18   750-046005   CACH5661         MPC5E 3D Q 2CGE+4XGE
  CPU
FPC 5            REV 35   750-028467   CAAR2623         MPC 3D 16x 10GE
 CPU
FPC 9            REV 30   750-044130   ABCF5773         MPC6E 3D
  CPU
FPC 10           REV 36   750-044130   ABCS8602         MPC6E 3D
  CPU
FPC 17           REV 28   750-044130   ABBZ3873         MPC6E 3D
  CPU
FPC 18           REV 39   750-045715   CACD1910         MPC5E 3D Q 24XGE+6XLGE
  CPU
FPC 19           REV 39   750-045715   CACD1908         MPC5E 3D Q 24XGE+6XLGE
  CPU
Fan Tray 0       REV 06   760-046960   ACAY0791         172mm FanTray - 6 Fans
Fan Tray 1       REV 06   760-046960   ACAY0788         172mm FanTray - 6 Fans
Fan Tray 2       REV 06   760-046960   ACAY0755         172mm FanTray - 6 Fans
Fan Tray 3       REV 06   760-046960   ACAY0441         172mm FanTray - 6 Fans
```

show chassis hardware extensive (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN120BADBAFJ      MX2020
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
                                 S/N:               JN120BADBAFJ
 Assembly ID:  0x0557            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: MX2020
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 30 42 41 44 42 41 46 4a 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 51   750-040240   ABAB9243          Lower Backplane
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-040240        S/N:               ABAB9243
 Assembly ID:  0x0b22            Assembly Version:  01.51
 Date:         05-30-2013        Assembly Flags:    0x00
 Version:      REV 51            CLEI Code:         IPMU710ARA
 ID: Lower Backplane             FRU Model Number:  CHAS-BP-MX2020-S
 Board Information Record:
  Address 0x00: ad 01 10 00 4c 96 14 72 30 08 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 22 01 33 52 45 56 20 35 31 00 00
```

```
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 34 30 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 39 32 34 33 00 1e 05 07
  Address 0x30: dd ff ff ff ad 01 10 00 4c 96 14 72 30 08 ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 55 37 31 30 41 52 41 43
  Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 32 30 2d 53 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff d3 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1       REV 04   711-032386   ABAB9399         Upper Backplane
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          711-032386       S/N:               ABAB9399
 Assembly ID:  0x0b23           Assembly Version:  01.04
 Date:         10-22-2012       Assembly Flags:    0x00
 Version:      REV 04
 ID: Upper Backplane
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 fe 0b 23 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 33 38 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 39 39 00 16 0a 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP 1            REV 05   711-032428   ACAJ2541         Upper Power Midplane
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          711-032428       S/N:               ACAJ2541
 Assembly ID:  0x045c           Assembly Version:  01.05
 Date:         04-26-2013       Assembly Flags:    0x00
 Version:      REV 05
 ID: Upper Power Midplane
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 5c 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 38 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 4a 32 35 34 31 00 1a 04 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP 0            REV 04   711-032426   ACAJ2194         Lower Power Midplane
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          711-032426       S/N:               ACAJ2194
 Assembly ID:  0x045d           Assembly Version:  01.04
 Date:         01-29-2013       Assembly Flags:    0x00
 Version:      REV 04
 ID: Lower Power Midplane
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 4a 32 31 39 34 00 1d 01 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```



```
FPM Board        REV 13   760-040242   ABCA8835          Front Panel Display
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          760-040242       S/N:               ABCA8835
 Assembly ID:  0x0b24           Assembly Version:  01.13
 Date:         04-13-2013       Assembly Flags:    0x00
 Version:      REV 13           CLEI Code:         IPMYAE5JRA
 ID: Front Panel Display        FRU Model Number:  MX2020-CRAFT-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 24 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 30 32 34 32 00 00
  Address 0x20: 53 2f 4e 20 41 42 43 41 38 38 33 35 00 0d 04 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 59 41 45 35 4a 52 41 4d
  Address 0x50: 58 32 30 32 30 2d 43 52 41 46 54 2d 53 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 95 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0            REV 01   740-050037   1EDB32403L5       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB32403L5
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         06-21-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 34 30 33 4c 35 00 00 15 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1            REV 01   740-050037   1EDB32403L3       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB32403L3
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         06-21-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 32 34 30 33 4c 33 00 00 15 06 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2            REV 01   740-050037   1EDB32403KM       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB32403KM
 Assembly ID:  0x0478           Assembly Version:  01.01
```

```
  Date:         06-21-2013        Assembly Flags:    0x00
  Version:      REV 01            CLEI Code:         IPUPAKRKAA
  ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
   Address 0x20: 31 45 44 42 33 32 34 30 33 4b 4d 00 00 15 06 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
   Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
   Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3            REV 01   740-050037   1EDB3130079       DC 52V Power Supply
Module
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          740-050037        S/N:               1EDB3130079
  Assembly ID:  0x0478            Assembly Version:  01.01
  Date:         05-16-2013        Assembly Flags:    0x00
  Version:      REV 01            CLEI Code:         IPUPAKRKAA
  ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
   Address 0x20: 31 45 44 42 33 31 33 30 30 37 39 00 00 10 05 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
   Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
   Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 4            REV 01   740-050037   1EDB3130077       DC 52V Power Supply
Module
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          740-050037        S/N:               1EDB3130077
  Assembly ID:  0x0478            Assembly Version:  01.01
  Date:         05-17-2013        Assembly Flags:    0x00
  Version:      REV 01            CLEI Code:         IPUPAKRKAA
  ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
   Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
   Address 0x20: 31 45 44 42 33 31 33 30 30 37 37 00 00 11 05 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
   Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
   Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 5            REV 01   740-050037   1EDB3130020       DC 52V Power Supply
Module
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          740-050037        S/N:               1EDB3130020
  Assembly ID:  0x0478            Assembly Version:  01.01
  Date:         05-16-2013        Assembly Flags:    0x00
  Version:      REV 01            CLEI Code:         IPUPAKRKAA
  ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
  Board Information Record:
```



```
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 31 33 30 30 32 30 00 00 10 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 6            REV 01   740-050037   1EDB313009S       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB313009S
 Assembly ID:  0x0478            Assembly Version:  01.01
 Date:         05-17-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 31 33 30 30 39 53 00 00 11 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7            REV 01   740-050037   1EDB313008E       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB313008E
 Assembly ID:  0x0478            Assembly Version:  01.01
 Date:         05-17-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 31 33 30 30 38 45 00 00 11 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 8            REV 01   740-050037   1EDB3130063       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB3130063
 Assembly ID:  0x0478            Assembly Version:  01.01
 Date:         05-17-2013        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
```

```
  Address 0x20: 31 45 44 42 33 31 33 30 30 36 33 00 00 11 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 12           REV 01   740-050037   1EDB3130026       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB3130026
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-16-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 31 33 30 30 32 36 00 00 10 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 13           REV 01   740-050037   1EDB3130074       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB3130074
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-17-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 31 33 30 30 37 34 00 00 11 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 14           REV 01   740-050037   1EDB313009D       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB313009D
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-17-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 33 31 33 30 30 39 44 00 00 11 05 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
```



```
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 15           REV 01   740-050037   1EDB3130024       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          740-050037       S/N:               1EDB3130024
 Assembly ID:  0x0478           Assembly Version:  01.01
 Date:         05-16-2013       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         IPUPAKRKAA
 ID: DC 52V Power Supply Module FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
...
```

show chassis hardware models (MX2020 Routers with MPC5EQ and MPC6E)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 51   750-040240   ABAB9243          CHAS-BP-MX2020-S
FPM Board        REV 13   760-040242   ABCA8835          MX2020-CRAFT-S
PSM 0            REV 01   740-050037   1EDB32403L5       MX2000-PSM-DC-S
PSM 1            REV 01   740-050037   1EDB32403L3       MX2000-PSM-DC-S
PSM 2            REV 01   740-050037   1EDB32403KM       MX2000-PSM-DC-S
PSM 3            REV 01   740-050037   1EDB3130079       MX2000-PSM-DC-S
PSM 4            REV 01   740-050037   1EDB3130077       MX2000-PSM-DC-S
PSM 5            REV 01   740-050037   1EDB3130020       MX2000-PSM-DC-S
PSM 6            REV 01   740-050037   1EDB313009S       MX2000-PSM-DC-S
PSM 7            REV 01   740-050037   1EDB313008E       MX2000-PSM-DC-S
PSM 8            REV 01   740-050037   1EDB3130063       MX2000-PSM-DC-S
PSM 12           REV 01   740-050037   1EDB3130026       MX2000-PSM-DC-S
PSM 13           REV 01   740-050037   1EDB3130074       MX2000-PSM-DC-S
PSM 14           REV 01   740-050037   1EDB313009D       MX2000-PSM-DC-S
PSM 15           REV 01   740-050037   1EDB3130024       MX2000-PSM-DC-S
PSM 16           REV 01   740-050037   1EDB3130054       MX2000-PSM-DC-S
PSM 17           REV 01   740-050037   1EDB3130080       MX2000-PSM-DC-S
PDM 0            REV 03   740-045234   1EGA3170144       MX2000-PDM-DC-S
PDM 1            REV 03   740-045234   1EGA3170158       MX2000-PDM-DC-S
PDM 2            REV 03   740-045234   1EGA3170182       MX2000-PDM-DC-S
PDM 3            REV 03   740-045234   1EGA3170207       MX2000-PDM-DC-S
Routing Engine 0 REV 02   740-041821   9009112112        RE-MX2000-1800X4-S
Routing Engine 1 REV 02   740-041821   9009112087        RE-MX2000-1800X4-S
CB 0             REV 23   750-040257   CABA2295          RE-MX2000-1800X4-S
CB 1             REV 23   750-040257   CABE8379          RE-MX2000-1800X4-S
SFB 0            REV 06   711-044466   ABCD5001          MX2000-SFB-S
SFB 1            REV 06   711-044466   ABCD5034          MX2000-SFB-S
SFB 2            REV 06   711-044466   ABCH3899          MX2000-SFB-S
SFB 3            REV 06   711-044466   ABCD5020          MX2000-SFB-S
SFB 4            REV 06   711-044466   ABCD4975          MX2000-SFB-S
SFB 5            REV 06   711-044466   ABCH3881          MX2000-SFB-S
SFB 6            REV 06   711-044466   ABCD5026          MX2000-SFB-S
SFB 7            REV 06   711-044466   ABCD5032          MX2000-SFB-S
FPC 0            REV 39   750-045715   CACD1902          PROTO-ASSEMBLY
FPC 1            REV 11   750-045372   CABK8112          MX-MPC3E-3D
FPC 2            REV 17   750-037355   CAAS5826          MPC4E-3D-2CGE-8XGE
FPC 3            REV 05   750-044444   CAAY9920          MX-MPC2E-3D-P
FPC 4            REV 18   750-046005   CACH5661          PROTO-ASSEMBLY
FPC 5            REV 35   750-028467   CAAR2623          MPC-3D-16XGE-SFPP
FPC 9            REV 30   750-044130   ABCF5773          PROTO-ASSEMBLY
```



```
FPC 10          REV 36   750-044130   ABCS8602          PROTO-ASSEMBLY
FPC 17          REV 28   750-044130   ABBZ3873          PROTO-ASSEMBLY
FPC 18          REV 39   750-045715   CACD1910          PROTO-ASSEMBLY
FPC 19          REV 39   750-045715   CACD1908          PROTO-ASSEMBLY
ADC 0           REV 17   750-043596   ABCD5378          MX2000-LC-ADAPTER
ADC 1           REV 17   750-043596   ABCD5465          MX2000-LC-ADAPTER
ADC 2           REV 17   750-043596   ABCD5431          MX2000-LC-ADAPTER
ADC 3           REV 17   750-043596   ABCD5356          MX2000-LC-ADAPTER
ADC 4           REV 02   750-043596   ZW1545            750-043596
ADC 5           REV 17   750-043596   ABCD5517          MX2000-LC-ADAPTER
ADC 18          REV 17   750-043596   ABCD5535          MX2000-LC-ADAPTER
ADC 19          REV 01   750-043596   ZV4127            750-043596
Fan Tray 0      REV 06   760-046960   ACAY0791          MX2000-FANTRAY-S
Fan Tray 1      REV 06   760-046960   ACAY0788          MX2000-FANTRAY-S
Fan Tray 2      REV 06   760-046960   ACAY0755          MX2000-FANTRAY-S
Fan Tray 3      REV 06   760-046960   ACAY0441          MX2000-FANTRAY-S
```

show chassis hardware clei-models (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item            Version  Part number  CLEI code         FRU model number
Midplane        REV 51   750-040240   IPMU710ARA        CHAS-BP-MX2020-S
FPM Board       REV 13   760-040242   IPMYAE5JRA        MX2020-CRAFT-S
PSM 0           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 1           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 2           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 3           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 4           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 5           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 6           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 7           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 8           REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 12          REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 13          REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 14          REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 15          REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 16          REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PSM 17          REV 01   740-050037   IPUPAKRKAA        MX2000-PSM-DC-S
PDM 0           REV 03   740-045234   IPUPAJSKAA        MX2000-PDM-DC-S
PDM 1           REV 03   740-045234   IPUPAJSKAA        MX2000-PDM-DC-S
PDM 2           REV 03   740-045234   IPUPAJSKAA        MX2000-PDM-DC-S
PDM 3           REV 03   740-045234   IPUPAJSKAA        MX2000-PDM-DC-S
CB 0            REV 23   750-040257   IPUCBA7CTA        RE-MX2000-1800X4-S
CB 1            REV 23   750-040257   IPUCBA7CTA        RE-MX2000-1800X4-S
SFB 0           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 1           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 2           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 3           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 4           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 5           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 6           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
SFB 7           REV 06   711-044466   IPUCBA6CAA        MX2000-SFB-S
FPC 0           REV 39   750-045715   PROTOXCLEI        PROTO-ASSEMBLY
FPC 1           REV 11   750-045372   COUIBBNBAA        MX-MPC3E-3D
FPC 2           REV 17   750-037355   IPU3A4DHAA        MPC4E-3D-2CGE-8XGE
FPC 3           REV 05   750-044444   COUIBBGBAA        MX-MPC2E-3D-P
  MIC 0         REV 28   750-028387   COUIA16BAA        MIC-3D-4XGE-XFP
FPC 4           REV 18   750-046005   PROTOXCLEI        PROTO-ASSEMBLY
FPC 5           REV 35   750-028467                     MPC-3D-16XGE-SFPP
FPC 9           REV 30   750-044130   PROTOXCLEI        PROTO-ASSEMBLY
```



```
  MIC 0          REV 05   750-049457   PROTOXCLEI        PROTO-ASSEMBLY
FPC 10           REV 36   750-044130   PROTOXCLEI        PROTO-ASSEMBLY
  MIC 0          REV 06   750-049979   PROTOXCLEI        PROTO-ASSEMBLY
  MIC 1          REV 12   750-050008   PROTOXCLEI        PROTO-ASSEMBLY
FPC 17           REV 28   750-044130   PROTOXCLEI        PROTO-ASSEMBLY
  MIC 1          REV 03   750-050008   PROTOXCLEI        PROTO-ASSEMBLY
FPC 18           REV 39   750-045715   PROTOXCLEI        PROTO-ASSEMBLY
FPC 19           REV 39   750-045715   PROTOXCLEI        PROTO-ASSEMBLY
ADC 0            REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
ADC 1            REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
ADC 2            REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
ADC 3            REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
ADC 4            REV 02   750-043596   PROTOXCLEI        750-043596
ADC 5            REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
ADC 18           REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
ADC 19           REV 01   750-043596   PROTOXCLEI        750-043596
Fan Tray 0       REV 06   760-046960   IPUCBA5CAA        MX2000-FANTRAY-S
Fan Tray 1       REV 06   760-046960   IPUCBA5CAA        MX2000-FANTRAY-S
Fan Tray 2       REV 06   760-046960   IPUCBA5CAA        MX2000-FANTRAY-S
Fan Tray 3       REV 06   760-046960   IPUCBA5CAA        MX2000-FANTRAY-S
```

show chassis hardware (MX Series routers with ATM MIC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN115736EAFC      MX240
Midplane         REV 07   760-021404   ABAA5038          MX240 Backplane
FPM Board        REV 03   760-021392   ABBA2758          Front Panel Display
PEM 0            Rev 01   740-022697   QCS0937C07K       PS 1.2-1.7kW; 100-240V
AC in
PEM 1            Rev 01   740-022697   QCS0939C04X       PS 1.2-1.7kW; 100-240V
AC in
PEM 2            Rev 01   740-022697   QCS0937C06B       PS 1.2-1.7kW; 100-240V
AC in
PEM 3            Rev 01   740-022697   QCS0937C07U       PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 12   740-013063   9009042291        RE-S-2000
Routing Engine 1 REV 12   740-013063   9009042266        RE-S-2000
CB 0             REV 06   710-021523   ABBC1435          MX SCB
CB 1             REV 06   710-021523   ABBC1497          MX SCB
FPC 2     REV 14    750-031088    YH8446           MPC Type 2 3D Q
  CPU            REV 06   711-030884   YH9612            MPC PMB 2G
  MIC 0
  MIC 1          REV 10   750-036132   ZP7062            2xOC12/8xOC3 CC-CE
    PIC 2                 BUILTIN      BUILTIN           2xOC12/8xOC3 CC-CE

  Xcvr 0              NON-JNPR     23393-00492       UNKNOWN
      Xcvr 1              NON-JNPR     23393-00500       UNKNOWN
      Xcvr 2              NON-JNPR     23393-00912       UNKNOWN
      Xcvr 3     REV 01   740-015638   22216-00575       Load SFP
      Xcvr 4     REV 01   740-015638   24145-00110       Load SFP
      Xcvr 5     REV 01   740-015638   24145-00016       Load SFP
      Xcvr 6     REV 01   740-015638   24145-00175       Load SFP
      Xcvr 7              NON-JNPR     23393-00627       UNKNOWN
  QXM 0          REV 05   711-028408   YF4681            MPC QXM
  QXM 1          REV 05   711-028408   YF4817            MPC QXM
Fan Tray 0       REV 01   710-021113   XL3645            MX240 Fan Tray
```


show chassis hardware (MX240, MX480, MX960 routers with Application Services Modular Line Card)

```
user@host>show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11D969BAFA      MX960
Midplane         REV 03   710-013698   ACAA2362          MX960 Backplane
FPM Board        REV 03   710-014974   ZR0639            Front Panel Display
PDM             Rev 03   740-013110   QCS152250SX      Power Distribution Module
PEM 0            Rev 10   740-013683   QCS1512718W       DC Power Entry Module
PEM 1            Rev 10   740-013683   QCS1512702Y       DC Power Entry Module
Routing Engine 0 REV 15   740-013063   9012024667        RE-S-2000
Routing Engine 1 REV 15   740-013063   9012024649        RE-S-2000
CB 0             REV 14   750-031391   ZJ7749            Enhanced MX SCB
CB 1             REV 14   750-031391   ZJ7750            Enhanced MX SCB
CB 2             REV 14   750-031391   ZY9233            Enhanced MX SCB
FPC 0            REV 17   750-031089   YR7434            MPC Type 2 3D
  CPU
FPC 1            REV 11   750-037207   ZW9727            AS-MCC
  CPU            REV 04   711-038173   ZW4817            AS-MCC-PMB
  MIC 0          REV 01   750-037214   ZH3764            AS-MSC
    PIC 0                 BUILTIN      BUILTIN           AS-MSC
  MIC 1          REV 01   711-028408   JZ9200            AS-MXC
    PIC 2                 BUILTIN      BUILTIN           AS-MXC
FPC 4            REV 30   750-028467   ABBN0232          MPC 3D 16x 10GE
  CPU
FPC 5            REV 04   750-037207   ZK9074            AS-MCC
  CPU
Fan Tray 0       REV 05   740-014971   VT5683            Fan Tray
Fan Tray 1       REV 05   740-014971   VT5684            Fan Tray
```

show chassis hardware extensive (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```
user@host> show chassis hardware extensive

ID: AS-MCC                      FRU Model Number:  750-037207
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
  Address 0x20: 53 2f 4e 20 5a 57 39 37 32 37 00 00 00 11 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 04   711-038173   ZW4817            AS-MCC-PMB
 Jedec Code:    0x7fb0           EEPROM Version:    0x02
 P/N:           711-038173       S/N:               ZW4817
 Assembly ID:   0x0b38           Assembly Version:  01.04
 Date:          12-30-2011       Assembly Flags:    0x00
 Version:       REV 04
 ID: AS-MCC-PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
```



```
  Address 0x20: 53 2f 4e 20 5a 57 34 38 31 37 00 00 00 1e 0c 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
  MIC 0          REV 01   750-037214   ZH3764           AS-MSC
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-037214       S/N:               ZH3764
 Assembly ID:  0x0a44           Assembly Version:  01.01
 Date:         07-04-2011       Assembly Flags:    0x00
 Version:      REV 01
 ID: AS-MSC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0a 44 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
  Address 0x20: 53 2f 4e 20 5a 48 33 37 36 34 00 00 00 04 07 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff f6 c0 03 e1 bc 00 00 00 00 00 00 00 00
    PIC 0                BUILTIN      BUILTIN          AS-MSC
FPC 4            REV 30   750-028467   ABBN0232         MPC 3D 16x 10GE
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
```

show chassis hardware (MX480 Router with MPC4E)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN10FF57BAFB      MX480
Midplane         REV 05   750-047849   Good              MX480 Midplane
FPM Board        REV 02   710-017254   KG2066            Front Panel Display
PEM 0            Rev 03   740-017330   QCS081590BJ       PS 1.2-1.7kW; 100-240V
AC in
PEM 1            Rev 03   740-017330   QCS0815908Z       PS 1.2-1.7kW; 100-240V
AC in
PEM 2            Rev 03   740-029970   QCS1001U001       PS 1.4-2.52kW; 90-264V
AC in
Routing Engine 0 REV 05   740-031116   9009089502        RE-S-1800x4
Routing Engine 1 REV 05   740-031116   9009089624        RE-S-1800x4
CB 0             REV 02   750-031391   YE8506            Enhanced MX SCB
CB 1             REV 14   750-031391   ZK8265            Enhanced MX SCB
FPC 2            REV 05   750-037358   ZT0638            MPC4E 3D 32XGE
  CPU            REV 07   711-035209   ZK3187            HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN           8X10GE SFPP
  PIC 1                   BUILTIN      BUILTIN           8X10GE SFPP
  PIC 2                   BUILTIN      BUILTIN           8X10GE SFPP
  PIC 3                   BUILTIN      BUILTIN           8X10GE SFPP
FPC 3            REV 06   750-037355   CAAB1144          MPC4E 3D 2CGE+8XGE
  CPU            REV 08   711-035209   CAAB1278          HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN           4x10GE SFPP
    Xcvr 0       REV 01   740-031980   B11E01439         SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11D05809         SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1X100GE CFP
    Xcvr 0                NON-JNPR     D5418             UNKNOWN
  PIC 2                   BUILTIN      BUILTIN           4x10GE SFPP
  PIC 3                   BUILTIN      BUILTIN           1X100GE CFP
    Xcvr 0                NON-JNPR     X12J00362         CFP-100G-SR10
```



```
FPC 4            REV 12.3.10 750-033205 YR9445          MPCE Type 3 3D
  CPU
Fan Tray                                                Enhanced Left Fan Tray
```

show chassis hardware (MX2020 Router with MPC4E)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11E188CAFJ      MX2020
Midplane         REV 04   711-032387   ABAC7474          Lower Backplane
Midplane 1       REV 04   711-032386   ABAC7408          Upper Backplane
PMP 1            REV 03   711-032428   ACAJ1137          Upper Power Midplane
PMP 0            REV 03   711-032426   ACAJ1016          Lower Power Midplane
FPM Board        REV 06   760-040242   ABBT8832          Front Panel Display
PSM 3            REV 0C   740-033727   VK00255           DC 52V Power Supply
Module
PSM 4            REV 0C   740-033727   VJ00148           DC 52V Power Supply
Module
PSM 5            REV 0C   740-033727   VK00207           DC 52V Power Supply
Module
PSM 6            REV 0C   740-033727   VK00319           DC 52V Power Supply
Module
PSM 7            REV 0C   740-033727   VK00264           DC 52V Power Supply
Module
PSM 8            REV 0B   740-033727   VG00025           DC 52V Power Supply
Module
PSM 13           REV 0C   740-033727   VK00274           DC 52V Power Supply
Module
PSM 14           REV 0C   740-033727   VJ00167           DC 52V Power Supply
Module
PSM 15           REV 0C   740-033727   VK00299           DC 52V Power Supply
Module
PSM 16           REV 0C   740-033727   VK00213           DC 52V Power Supply
Module
PSM 17           REV 0C   740-033727   VK00253           DC 52V Power Supply
Module
PDM 0            REV 0B   740-038109   VJ00040           DC Power Dist Module
PDM 2            REV 0B   740-038109   VJ00025           DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009089735        RE-S-1800x4
Routing Engine 1 REV 02   740-041821   9009089731        RE-S-1800x4
CB 0             REV 04   750-040257   ZT2846            Control Board
CB 1             REV 04   750-040257   ZT2877            Control Board
SPMB 0           REV 01   711-041855   ZS2282            PMB Board
SPMB 1           REV 01   711-041855   ZS2261            PMB Board
SFB 0            REV 07   711-032385   ZZ2582            Switch Fabric Board
SFB 1            REV 04   711-032385   ZV4229            Switch Fabric Board
SFB 2            REV 07   711-032385   CAAB4902          Switch Fabric Board
SFB 3            REV 07   711-032385   CAAB4891          Switch Fabric Board
SFB 4            REV 07   711-032385   CAAB4883          Switch Fabric Board
SFB 5            REV 07   711-032385   CAAB4889          Switch Fabric Board
SFB 6            REV 06   711-032385   ZV1818            Switch Fabric Board
SFB 7            REV 07   711-032385   CAAB4897          Switch Fabric Board
FPC 0            REV 34   750-031090   ZT9799            MPC Type 2 3D EQ
  CPU            REV 06   711-030884   ZS1122            MPC PMB 2G
  MIC 0          REV 11   750-033535   CAAD7674          MIC-3D-1OC192-XFP
    PIC 0                 BUILTIN      BUILTIN           MIC-3D-1OC192-XFP
      Xcvr 0     REV 01   740-014279   753019A00404      XFP-OC192-SR
  MIC 1          REV 14   750-031967   ZM6103            MIC-3D-8OC3OC12-4OC48
    PIC 2                 BUILTIN      BUILTIN           MIC-3D-8OC3OC12-4OC48
      Xcvr 0     REV 01   740-011615   PEF1AZP           SFP-IR
```

```
      Xcvr 1     REV 01   740-011615   PEF1AZN          SFP-IR
      Xcvr 2     REV 01   740-021308   ANA0N8S          SFP+-10G-SR
  QXM 0          REV 06   711-028408   ZT9339           MPC QXM
  QXM 1          REV 06   711-028408   ZT9237           MPC QXM
FPC 9            REV 34   750-031090   ZT9770           MPC Type 2 3D EQ
  CPU            REV 06   711-030884   ZS1302           MPC PMB 2G
  MIC 0          REV 24   750-028387   YJ3950           3D 4x 10GE  XFP
    PIC 0                 BUILTIN      BUILTIN          2x 10GE  XFP
      Xcvr 0              NON-JNPR     T09M52516        XFP-10G-SR
      Xcvr 1              NON-JNPR     CA49BK095        XFP-10G-SR
    PIC 1                 BUILTIN      BUILTIN          2x 10GE  XFP
      Xcvr 0     REV 02   740-014289   C834XU01T        XFP-10G-SR
      Xcvr 1              NON-JNPR     T09M52515        XFP-10G-SR
  MIC 1          REV 11   750-033535   CAAD7681         MIC-3D-1OC192-XFP
    PIC 2                 BUILTIN      BUILTIN          MIC-3D-1OC192-XFP
      Xcvr 0     REV 01   740-014279   KBQ02BE          XFP-OC192-SR
  QXM 0          REV 06   711-028408   ZT9151           MPC QXM
  QXM 1          REV 06   711-028408   ZT9116           MPC QXM
FPC 10           REV 27   750-033205   ZL6215           MPCE Type 3 3D
  CPU            REV 07   711-035209   ZK9038           HMPC PMB 2G
  MIC 0          REV 18   750-028380   YG6885           3D 2x 10GE XFP
    PIC 0                 BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0     REV 01   740-014289   C706XU0AG        XFP-10G-SR
    PIC 1                 BUILTIN      BUILTIN          1x 10GE XFP
      Xcvr 0     REV 02   740-014289   T08L84366        XFP-10G-SR
FPC 14           REV 09   750-037355   CAAF1534         MPC4E 3D 2CGE+8XGE
  CPU            REV 08   711-035209   CAAB9879         HMPC PMB 2G
  PIC 0                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-021308   21T511100436     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AHP0GPM          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   123363A00032     SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   19T511100477     SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          1X100GE CFP
    Xcvr 0                NON-JNPR     X12J00260        CFP-100G-SR10
  PIC 2                   BUILTIN      BUILTIN          4x10GE SFPP
    Xcvr 0       REV 01   740-021308   21T511104086     SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   21T511104627     SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   21T511104644     SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          1X100GE CFP
FPC 19           REV 32   750-028467   ZR2008           MPC 3D 16x 10GE
  CPU            REV 10   711-029089   ZT6933           AMPC PMB
  PIC 0                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   19T511100291     SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AMH02VE          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   23T511102128     SFP+-10G-SR
  PIC 2                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-021308   AMS15PP          SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   123363A00716     SFP+-10G-SR
ADC 0            REV 05   750-043596   CAAC2072         Adapter Card
ADC 9            REV 01   750-043596   ZV4111           Adapter Card
ADC 10           REV 05   750-043596   CAAC2058         Adapter Card
ADC 14           REV 02   750-043596   ZW1561           Adapter Card
ADC 19           REV 01   750-043596   ZV4127           Adapter Card
Fan Tray 0       REV 03   760-046960   ACAY0124         172mm FanTray - 6 Fans
Fan Tray 1       REV 2A   760-046960   ACAY0022         172mm FanTray - 6 Fans
Fan Tray 2       REV 2A   760-046960   ACAY0023         172mm FanTray - 6 Fans
Fan Tray 3       REV 2A   760-046960   ACAY0025         172mm FanTray - 6 Fans
```

show chassis hardware (MX5, MX10, MX40, MX80, MX240, MX480, and MX960 Routers with Enhanced 20-Port Gigabit Ethernet MIC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                F3434             MX80-P
Midplane         REV 01   711-044315   ZK2681            MX80-P
PEM 0            Rev 04   740-028288   VE05267           AC Power Entry Module
PEM 1            Rev 04   740-028288   VE05270           AC Power Entry Module
Routing Engine            BUILTIN      BUILTIN           Routing Engine
TFEB 0                    BUILTIN      BUILTIN           Forwarding Engine
Processor
  QXM 0          REV 05   711-028408   ZK0952            MPC QXM
FPC 0                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN           4x 10GE XFP
    PIC 0                 BUILTIN      BUILTIN           4x 10GE XFP
FPC 1                     BUILTIN      BUILTIN           MPC BUILTIN
  MIC 0          REV 02   750-049846   CAAV2153          3D 20x 1GE(LAN)-E,SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) -E  SFP
      Xcvr 0     REV 01   740-011613   AM0816S9B81       SFP-SX
      Xcvr 1     REV 02   740-011613   AM0925SBLK7       SFP-SX
      Xcvr 2     REV 01   740-011613   UAQ0005           SFP-SX
      Xcvr 3     REV 01   740-011613   UAQ000C           SFP-SX
      Xcvr 4     REV 01   740-011613   P9F195E           SFP-SX
      Xcvr 5     REV 01   740-011613   UAQ0003           SFP-SX
      Xcvr 6     REV 01   740-031851   AM1041SU1LD       SFP-SX
      Xcvr 8     REV 02   740-013111   B101501           SFP-T
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) -E  SFP
      Xcvr 0     REV 01   740-011613   PFM1ML7           SFP-SX
      Xcvr 4     REV 01   740-011613   PE729P6           SFP-SX
      Xcvr 6     REV 02   740-011613   AM1014SGC84       SFP-SX
      Xcvr 9     REV 01   740-011613   AM0812S8UK3       SFP-SX
  MIC 1          REV 26   750-028392   ZY0187            3D 20x 1GE(LAN) SFP
    PIC 2                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-011613   P9F1AN9           SFP-SX
      Xcvr 5     REV 02   740-011613   AM1003SFUF4       SFP-SX
      Xcvr 9     REV 01   740-031851   AM1041SU1LM       SFP-SX
    PIC 3                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 4     REV 01   740-011613   PAJ4MYT           SFP-SX
      Xcvr 7     +        NON-JNPR     XG32A024          SFP-SX
      Xcvr 8              NON-JNPR     PFR0V6J           SFP-SX
      Xcvr 9     REV 01   740-031851   AM1041SU02U       SFP-SX
Fan Tray
```

show chassis hardware models (MX5, MX10, MX40, MX80, MX240, MX480, and MX960 Routers with Enhanced 20-Port Gigabit Ethernet MIC)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
PEM 0            Rev 04   740-028288   VE05267           PWR-MX80-AC-S
PEM 1            Rev 04   740-028288   VE05270           PWR-MX80-AC-S
Routing Engine            BUILTIN      BUILTIN
TFEB 0                    BUILTIN      BUILTIN
FPC 0                     BUILTIN      BUILTIN
FPC 1                     BUILTIN      BUILTIN
  MIC 0          REV 02   750-049846   CAAV2153          MIC-3D-20GE-SFP-E
  MIC 1          REV 26   750-028392   ZY0187            MIC-3D-20GE-SFP
Fan Tray                                                 FANTRAY-MX80-S
```

show chassis hardware (MX2008 Router)

```
user@host>show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1259E1CAFL      MX2008
Midplane         REV 47   750-044636   ABAD1739          Lower Backplane
PMP              REV 01   711-051406   ACVD0738          Power Midplane
FPM Board        REV 02   760-068193   ABDG7408          Front Panel Display
PSM 1            REV 06   740-050037   1EDB61200R8       DC 52V Power Supply
Module
PSM 2            REV 06   740-050037   1EDB61200WA       DC 52V Power Supply
Module
PSM 3            REV 06   740-050037   1EDB61200NY       DC 52V Power Supply
Module
PSM 4            REV 06   740-050037   1EDB61200N2       DC 52V Power Supply
Module
PSM 5            REV 06   740-050037   1EDB61200RN       DC 52V Power Supply
Module
PSM 6            REV 06   740-050037   1EDB61200RF       DC 52V Power Supply
Module
PSM 7            REV 06   740-050037   1EDB61200R7       DC 52V Power Supply
Module
PDM 0            REV 01   740-060189   1EFF5250143       DC PDM Optimized
PDM 1            REV 01   740-060189   1EFF5250074       DC PDM Optimized
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x8
Routing Engine 1          BUILTIN      BUILTIN           RE-S-2X00x8
CB 0             REV 01   750-067373   ABDJ0047          Control Board
CB 1             REV 03   750-067373   ABDH3016          Control Board
SFB 0            REV 08   750-067371   ABDK7180          Switch Fabric Board
SFB 1            REV 08   750-067371   ABDK7024          Switch Fabric Board
SFB 2            REV 08   750-067371   ABDK7188          Switch Fabric Board
SFB 3            REV 08   750-067371   ABDK7143          Switch Fabric Board
SFB 4            REV 08   750-067371   ABDK7030          Switch Fabric Board
SFB 5            REV 08   750-067371   ABDK7146          Switch Fabric Board
SFB 6            REV 08   750-067371   ABDK7203          Switch Fabric Board
SFB 7            REV 08   750-067371   ABDK7238          Switch Fabric Board
FPC 0            REV 36   750-044130   ABCS8607          MPC6E 3D
  CPU            REV 09   711-045719   ABCS8776          RMPC PMB
  MIC 0          REV 21   750-050008   ABCT5920          4X100GE CXP
    PIC 0                 BUILTIN      BUILTIN           4X100GE CXP
  XLM 0          REV 07.2.00 711-046638 ABCK3488         MPC6E XL
  XLM 1          REV 07.2.00 711-046638 ABCK5482         MPC6E XL
FPC 1            REV 22   750-063414   CAFJ3026          MPC9E 3D
  CPU            REV 16   750-057177   CAFF9332          SMPC PMB
FPC 7            REV 08   750-038492   ZX4080            MPCE Type 2 3D EQ
  CPU            REV 03   711-038484   ZX3665            MPCE PMB 2G
  MIC 0          REV 05   750-037128   ZR4031            1xCOC12/4xCOC3 CH-CE
    PIC 0                 BUILTIN      BUILTIN           1xCOC12/4xCOC3 CH-CE
  MIC 1          REV 23   750-032479   CADE8614          MIC-3D-8DS3-E3
    PIC 2                 BUILTIN      BUILTIN           MIC-3D-8DS3-E3
  QXM 0          REV 06   711-028408   ZW8299            MPC QXM
  QXM 1          REV 06   711-028408   ZY0609            MPC QXM
ADC 7            REV 17   750-043596   ABCA0990          Adapter Card
Fan Tray 0       REV 01   760-052467   ACAY6190          172mm FanTray - 6 Fans
Fan Tray 1       REV 01   760-052467   ACAY6414          172mm FanTray - 6 Fans
```

show chassis hardware detail (MX2008 Router)

```
user@host>show chassis hardware detail
```

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1259E1CAFL      MX2008
Midplane         REV 47   750-044636   ABAD1739          Lower Backplane
PMP              REV 01   711-051406   ACVD0738          Power Midplane
FPM Board        REV 02   760-068193   ABDG7408          Front Panel Display
PSM 1            REV 06   740-050037   1EDB61200R8       DC 52V Power Supply
Module
PSM 2            REV 06   740-050037   1EDB61200WA       DC 52V Power Supply
Module
PSM 3            REV 06   740-050037   1EDB61200NY       DC 52V Power Supply
Module
PSM 4            REV 06   740-050037   1EDB61200N2       DC 52V Power Supply
Module
PSM 5            REV 06   740-050037   1EDB61200RN       DC 52V Power Supply
Module
PSM 6            REV 06   740-050037   1EDB61200RF       DC 52V Power Supply
Module
PSM 7            REV 06   740-050037   1EDB61200R7       DC 52V Power Supply
Module
PDM 0            REV 01   740-060189   1EFF5250143       DC PDM Optimized
PDM 1            REV 01   740-060189   1EFF5250074       DC PDM Optimized
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x8
  vtbd0 15361 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    511 MB  QEMU HARDDISK        QM00002           Emulated IDE Disk
  usb0 (addr 1)  XHCI root HUB 0       0x8086            uhub0
Routing Engine 1          BUILTIN      BUILTIN           RE-S-2X00x8
  vtbd0 15361 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    511 MB  QEMU HARDDISK        QM00002           Emulated IDE Disk
  usb0 (addr 1)  XHCI root HUB 0       0x8086            uhub0
CB 0             REV 01   750-067373   ABDJ0047          Control Board
CB 1             REV 03   750-067373   ABDH3016          Control Board
SFB 0            REV 08   750-067371   ABDK7180          Switch Fabric Board
SFB 1            REV 08   750-067371   ABDK7024          Switch Fabric Board
SFB 2            REV 08   750-067371   ABDK7188          Switch Fabric Board
SFB 3            REV 08   750-067371   ABDK7143          Switch Fabric Board
SFB 4            REV 08   750-067371   ABDK7030          Switch Fabric Board
SFB 5            REV 08   750-067371   ABDK7146          Switch Fabric Board
SFB 6            REV 08   750-067371   ABDK7203          Switch Fabric Board
SFB 7            REV 08   750-067371   ABDK7238          Switch Fabric Board
FPC 0            REV 36   750-044130   ABCS8607          MPC6E 3D
  CPU            REV 09   711-045719   ABCS8776          RMPC PMB
  MIC 0          REV 21   750-050008   ABCT5920          4X100GE CXP
    PIC 0                 BUILTIN      BUILTIN           4X100GE CXP
  XLM 0          REV 07.2.00 711-046638 ABCK3488         MPC6E XL
  XLM 1          REV 07.2.00 711-046638 ABCK5482         MPC6E XL
FPC 1            REV 22   750-063414   CAFJ3026          MPC9E 3D
  CPU            REV 16   750-057177   CAFF9332          SMPC PMB
FPC 7            REV 08   750-038492   ZX4080            MPCE Type 2 3D EQ
  CPU            REV 03   711-038484   ZX3665            MPCE PMB 2G
  MIC 0          REV 05   750-037128   ZR4031            1xCOC12/4xCOC3 CH-CE
    PIC 0                 BUILTIN      BUILTIN           1xCOC12/4xCOC3 CH-CE
  MIC 1          REV 23   750-032479   CADE8614          MIC-3D-8DS3-E3
    PIC 2                 BUILTIN      BUILTIN           MIC-3D-8DS3-E3
  QXM 0          REV 06   711-028408   ZW8299            MPC QXM
  QXM 1          REV 06   711-028408   ZY0609            MPC QXM
ADC 7            REV 17   750-043596   ABCA0990          Adapter Card
Fan Tray 0       REV 01   760-052467   ACAY6190          172mm FanTray - 6 Fans
Fan Tray 1       REV 01   760-052467   ACAY6414          172mm FanTray - 6 Fans
```


show chassis hardware extensive (MX2008 Router)

```
user@host>show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1259E1CAFL      MX2008
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
                                S/N:               JN1259E1CAFL
 Assembly ID:  0x0557           Assembly Version:  00.00
 Date:         00-00-0000       Assembly Flags:    0x00
 ID: MX2008
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 32 35 39 45 31 43 41 46 4c 00 00 00 00
  Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 47   750-044636   ABAD1739          Lower Backplane
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-044636       S/N:               ABAD1739
 Assembly ID:  0x0b66           Assembly Version:  01.47
 Date:         06-08-2016       Assembly Flags:    0x00
 Version:      REV 47           CLEI Code:         IPMU810ARB
 ID: Lower Backplane            FRU Model Number:  CHAS-BP-MX2010-S
 Board Information Record:
  Address 0x00: ad 01 08 00 f4 cc 55 3e 35 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 66 01 2f 52 45 56 20 34 37 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 44 31 37 33 39 00 08 06 07
  Address 0x30: e0 ff ff ff ad 01 08 00 f4 cc 55 3e 35 00 ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 55 38 31 30 41 52 42 43
  Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 31 30 2d 53 00
  Address 0x60: 00 00 00 00 00 00 42 43 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 18 ff ff ff ff ff ff ff ff ff ff ff ff
PMP              REV 01   711-051406   ACVD0738          Power Midplane
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          711-051406       S/N:               ACVD0738
 Assembly ID:  0x045d           Assembly Version:  01.01
 Date:         06-06-2016       Assembly Flags:    0x00
 Version:      REV 01
 ID: Power Midplane
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 5d 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 35 31 34 30 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 56 44 30 37 33 38 00 06 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board        REV 02   760-068193   ABDG7408          Front Panel Display
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
```

```
 P/N:           760-068193        S/N:               ABDG7408
 Assembly ID:  0x0cac            Assembly Version:  01.02
 Date:          06-06-2016        Assembly Flags:    0x00
 Version:       REV 02            CLEI Code:         PROTOXCLEI
 ID: Front Panel Display         FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c ac 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 36 38 31 39 33 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 47 37 34 30 38 00 06 06 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 1            REV 06   740-050037   1EDB61200R8       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:           740-050037        S/N:               1EDB61200R8
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:          03-16-2016        Assembly Flags:    0x00
 Version:       REV 06            CLEI Code:         IPUPAPDKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 52 38 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2            REV 06   740-050037   1EDB61200WA       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:           740-050037        S/N:               1EDB61200WA
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:          03-16-2016        Assembly Flags:    0x00
 Version:       REV 06            CLEI Code:         IPUPAPDKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 57 41 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3            REV 06   740-050037   1EDB61200NY       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:           740-050037        S/N:               1EDB61200NY
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:          03-16-2016        Assembly Flags:    0x00
 Version:       REV 06            CLEI Code:         IPUPAPDKAA
```



```
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 4e 59 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PSM 4            REV 06   740-050037   1EDB61200N2       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB61200N2
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:         03-16-2016        Assembly Flags:    0x00
 Version:      REV 06            CLEI Code:         IPUPAPDKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 4e 32 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PSM 5            REV 06   740-050037   1EDB61200RN       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB61200RN
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:         03-16-2016        Assembly Flags:    0x00
 Version:      REV 06            CLEI Code:         IPUPAPDKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 52 4e 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PSM 6            REV 06   740-050037   1EDB61200RF       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB61200RF
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:         03-16-2016        Assembly Flags:    0x00
 Version:      REV 06            CLEI Code:         IPUPAPDKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
```

```
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 52 46 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7            REV 06   740-050037   1EDB61200R7       DC 52V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-050037        S/N:               1EDB61200R7
 Assembly ID:  0x0478            Assembly Version:  01.06
 Date:         03-16-2016        Assembly Flags:    0x00
 Version:      REV 06            CLEI Code:         IPUPAPDKAA
 ID: DC 52V Power Supply Module  FRU Model Number:  MX2000-PSM-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
  Address 0x20: 31 45 44 42 36 31 32 30 30 52 37 00 00 10 03 07
  Address 0x30: e0 72 75 ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 50 44 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
  Address 0x60: 00 00 00 00 00 00 31 30 36 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 26 00 00 00 00 00 00 00 00 00 00 00 00
PDM 0            REV 01   740-060189   1EFF5250143       DC PDM Optimized
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-060189        S/N:               1EFF5250143
 Assembly ID:  0x0495            Assembly Version:  01.01
 Date:         07-21-2015        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAN1KAA
 ID: DC PDM Optimized            FRU Model Number:  MX2K-PDM-OP-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 95 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 36 30 31 38 39 00 00
  Address 0x20: 31 45 46 46 35 32 35 30 31 34 33 00 00 15 07 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4e 31 4b 41 41 4d
  Address 0x50: 58 32 4b 2d 50 44 4d 2d 4f 50 2d 44 43 2d 53 00
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 84 00 00 00 00 00 00 00 00 00 00 00 00
PDM 1            REV 01   740-060189   1EFF5250074       DC PDM Optimized
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-060189        S/N:               1EFF5250074
 Assembly ID:  0x0495            Assembly Version:  01.01
 Date:         07-21-2015        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUPAN1KAA
 ID: DC PDM Optimized            FRU Model Number:  MX2K-PDM-OP-DC-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 95 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 36 30 31 38 39 00 00
  Address 0x20: 31 45 46 46 35 32 35 30 30 37 34 00 00 15 07 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4e 31 4b 41 41 4d
  Address 0x50: 58 32 4b 2d 50 44 4d 2d 4f 50 2d 44 43 2d 53 00
```

```
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 84 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x8
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0c10            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: RE-S-2X00x8
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0c 10 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 00 00 00 00
  Address 0x20: 42 55 49 4c 54 49 4e 00 00 00 00 00 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  vtbd0 15361 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    511 MB  QEMU HARDDISK        QM00002           Emulated IDE Disk
  usb0 (addr 1)  XHCI root HUB 0       0x8086            uhub0
Routing Engine 1          BUILTIN      BUILTIN           RE-S-2X00x8
 Jedec Code:   0x0000            EEPROM Version:    0x00
 P/N:          BUILTIN           S/N:               BUILTIN
 Assembly ID:  0x0c10            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: RE-S-2X00x8
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0c 10 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 00 00 00 00
  Address 0x20: 42 55 49 4c 54 49 4e 00 00 00 00 00 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  vtbd0 15361 MB                                         Virtio Block Disk
  vtbd1 15360 MB                                         Virtio Block Disk
  ada0    511 MB  QEMU HARDDISK        QM00002           Emulated IDE Disk
  usb0 (addr 1)  XHCI root HUB 0       0x8086            uhub0
CB 0             REV 01   750-067373   ABDJ0047          Control Board
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-067373        S/N:               ABDJ0047
 Assembly ID:  0x0c96            Assembly Version:  01.01
 Date:         06-21-2016        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         PROTOXCLEI
 ID: Control Board               FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ad 01 00 20 28 8a 1c 6d c4 7e ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 96 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 33 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 4a 30 30 34 37 00 15 06 07
  Address 0x30: e0 ff ff ff ad 01 00 20 28 8a 1c 6d c4 7e ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
```

```
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
CB 1             REV 03   750-067373   ABDH3016         Control Board
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-067373       S/N:               ABDH3016
 Assembly ID:  0x0c96           Assembly Version:  01.03
 Date:         05-07-2016       Assembly Flags:    0x00
 Version:      REV 03           CLEI Code:         PROTOXCLEI
 ID: Control Board              FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ad 01 00 20 f4 cc 55 35 71 a0 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 96 01 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 33 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 48 33 30 31 36 00 07 05 07
  Address 0x30: e0 ff ff ff ad 01 00 20 f4 cc 55 35 71 a0 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
SFB 0            REV 08   750-067371   ABDK7180         Switch Fabric Board
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-067371       S/N:               ABDK7180
 Assembly ID:  0x0c97           Assembly Version:  01.08
 Date:         09-27-2016       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         PROTOXCLEI
 ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 4b 37 31 38 30 00 1b 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 00 00 48 00
SFB 1            REV 08   750-067371   ABDK7024         Switch Fabric Board
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-067371       S/N:               ABDK7024
 Assembly ID:  0x0c97           Assembly Version:  01.08
 Date:         09-27-2016       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         PROTOXCLEI
 ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 4b 37 30 32 34 00 1b 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 01 00 48 00
SFB 2            REV 08   750-067371   ABDK7188         Switch Fabric Board
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-067371       S/N:               ABDK7188
 Assembly ID:  0x0c97           Assembly Version:  01.08
 Date:         09-28-2016       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         PROTOXCLEI
```



```
  ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
   Address 0x20: 53 2f 4e 20 41 42 44 4b 37 31 38 38 00 1c 09 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
   Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
   Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 02 00 48 00
SFB 3            REV 08   750-067371   ABDK7143         Switch Fabric Board
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-067371       S/N:               ABDK7143
  Assembly ID:  0x0c97           Assembly Version:  01.08
  Date:         09-27-2016       Assembly Flags:    0x00
  Version:      REV 08           CLEI Code:         PROTOXCLEI
  ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
   Address 0x20: 53 2f 4e 20 41 42 44 4b 37 31 34 33 00 1b 09 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
   Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
   Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 03 00 48 00
SFB 4            REV 08   750-067371   ABDK7030         Switch Fabric Board
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-067371       S/N:               ABDK7030
  Assembly ID:  0x0c97           Assembly Version:  01.08
  Date:         09-24-2016       Assembly Flags:    0x00
  Version:      REV 08           CLEI Code:         PROTOXCLEI
  ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
   Address 0x20: 53 2f 4e 20 41 42 44 4b 37 30 33 30 00 18 09 07
   Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
   Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
   Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 04 00 48 00
SFB 5            REV 08   750-067371   ABDK7146         Switch Fabric Board
  Jedec Code:   0x7fb0           EEPROM Version:    0x02
  P/N:          750-067371       S/N:               ABDK7146
  Assembly ID:  0x0c97           Assembly Version:  01.08
  Date:         09-27-2016       Assembly Flags:    0x00
  Version:      REV 08           CLEI Code:         PROTOXCLEI
  ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
   Address 0x20: 53 2f 4e 20 41 42 44 4b 37 31 34 36 00 1b 09 07
```

```
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 05 00 48 00
SFB 6            REV 08   750-067371   ABDK7203         Switch Fabric Board
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-067371       S/N:               ABDK7203
 Assembly ID:  0x0c97           Assembly Version:  01.08
 Date:         09-28-2016       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         PROTOXCLEI
 ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 4b 37 32 30 33 00 1c 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 06 00 48 00
SFB 7            REV 08   750-067371   ABDK7238         Switch Fabric Board
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-067371       S/N:               ABDK7238
 Assembly ID:  0x0c97           Assembly Version:  01.08
 Date:         09-27-2016       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         PROTOXCLEI
 ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0c 97 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 37 33 37 31 00 00
  Address 0x20: 53 2f 4e 20 41 42 44 4b 37 32 33 38 00 1b 09 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 07 00 00 00 00 00 00 00 07 00 48 00
FPC 0            REV 36   750-044130   ABCS8607         MPC6E 3D
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-044130       S/N:               ABCS8607
 Assembly ID:  0x0b86           Assembly Version:  01.36
 Date:         10-29-2013       Assembly Flags:    0x00
 Version:      REV 36           CLEI Code:         PROTOXCLEI
 ID: MPC6E 3D                   FRU Model Number:  PROTO-ASSEMBLY
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 86 01 24 52 45 56 20 33 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 31 33 30 00 00
  Address 0x20: 53 2f 4e 20 41 42 43 53 38 36 30 37 00 1d 0a 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 09   711-045719   ABCS8776         RMPC PMB
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
```



```
P/N:            711-045719        S/N:              ABCS8776
Assembly ID:    0x0b85            Assembly Version:  01.09
Date:           10-24-2013        Assembly Flags:    0x00
Version:        REV 09
ID: RMPC PMB
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 85 01 09 52 45 56 20 30 39 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 34 35 37 31 39 00 00
 Address 0x20: 53 2f 4e 20 41 42 43 53 38 37 37 36 00 18 0a 07
 Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 00 00 00 00 16 47 1f b0 00 00 00 00
 MIC 0          REV 21   750-050008   ABCT5920          4X100GE CXP
Jedec Code:     0x7fb0            EEPROM Version:    0x02
P/N:            750-050008        S/N:              ABCT5920
Assembly ID:    0x0a83            Assembly Version:  01.21
Date:           09-29-2014        Assembly Flags:    0x00
Version:        REV 21            CLEI Code:         IP9IATYDAA
ID: 4X100GE CXP                   FRU Model Number:  MIC6-100G-CXP
Board Information Record:
 Address 0x00: 12 01 07 02 03 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0a 83 01 15 52 45 56 20 32 31 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 35 30 30 30 38 00 00
 Address 0x20: 53 2f 4e 20 41 42 43 54 35 39 32 30 00 1d 09 07
 Address 0x30: de ff ff ff 12 01 07 02 03 ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 49 50 39 49 41 54 59 44 41 41 4d
 Address 0x50: 49 43 36 2d 31 30 30 47 2d 43 58 50 00 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 41 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff 74 00 00 00 00 10 09 73 3c c0 02 70 3c
   PIC 0                 BUILTIN      BUILTIN           4X100GE CXP
 XLM 0          REV 07.2.00 711-046638 ABCK3488         MPC6E XL
Jedec Code:     0x7fb0            EEPROM Version:    0x02
P/N:            711-046638        S/N:              ABCK3488
Assembly ID:    0x0b88            Assembly Version:  01.07
Date:           11-11-2013        Assembly Flags:    0x00
Version:        REV 07.2.00
ID: MPC6E XL
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 88 01 07 52 45 56 20 30 37 2e 32
 Address 0x10: 2e 30 30 00 37 31 31 2d 30 34 36 36 33 38 00 00
 Address 0x20: 53 2f 4e 20 41 42 43 4b 33 34 38 38 00 0b 0b 07
 Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 00 00 00 00 00 00 00 00 00 00 00 00
 XLM 1          REV 07.2.00 711-046638 ABCK5482         MPC6E XL
Jedec Code:     0x7fb0            EEPROM Version:    0x02
P/N:            711-046638        S/N:              ABCK5482
Assembly ID:    0x0b88            Assembly Version:  01.07
Date:           10-21-2013        Assembly Flags:    0x00
Version:        REV 07.2.00
ID: MPC6E XL
Board Information Record:
```



```
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 88 01 07 52 45 56 20 30 37 2e 32
  Address 0x10: 2e 30 30 00 37 31 31 2d 30 34 36 36 33 38 00 00
  Address 0x20: 53 2f 4e 20 41 42 43 4b 35 34 38 32 00 15 0a 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 00 00 00 00 00 00 00 00 00 00 00 00
FPC 1            REV 22   750-063414   CAFJ3026         MPC9E 3D
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-063414       S/N:               CAFJ3026
 Assembly ID:  0x0c43           Assembly Version:  01.22
 Date:         03-28-2016       Assembly Flags:    0x00
 Version:      REV 22           CLEI Code:         IPUCBMUCAA
 ID: MPC9E 3D                   FRU Model Number:  MX2K-MPC9E
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0c 43 01 16 52 45 56 20 32 32 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 36 33 34 31 34 00 00
  Address 0x20: 53 2f 4e 20 43 41 46 4a 33 30 32 36 00 1c 03 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 43 42 4d 55 43 41 41 4d
  Address 0x50: 58 32 4b 2d 4d 50 43 39 45 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 16   750-057177   CAFF9332         SMPC PMB
 Jedec Code:   0x7fb0           EEPROM Version:    0x01
 P/N:          750-057177       S/N:               CAFF9332
 Assembly ID:  0x0c22           Assembly Version:  01.16
 Date:         03-20-2016       Assembly Flags:    0x00
 Version:      REV 16
 ID: SMPC PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0c 22 01 10 52 45 56 20 31 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 35 37 31 37 37 00 00
  Address 0x20: 53 2f 4e 20 43 41 46 46 39 33 33 32 00 14 03 07
  Address 0x30: e0 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff 00 00 00 00 38 f9 0d e0 4f d1 4b 08
FPC 7            REV 08   750-038492   ZX4080           MPCE Type 2 3D EQ
 Jedec Code:   0x7fb0           EEPROM Version:    0x02
 P/N:          750-038492       S/N:               ZX4080
 Assembly ID:  0x0b35           Assembly Version:  01.08
 Date:         02-03-2012       Assembly Flags:    0x00
 Version:      REV 08           CLEI Code:         COUIBA5BAA
 ID: MPCE Type 2 3D EQ          FRU Model Number:  MX-MPC2E-3D-EQ
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 35 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 38 34 39 32 00 00
  Address 0x20: 53 2f 4e 20 5a 58 34 30 38 30 00 00 00 03 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 43 4f 55 49 42 41 35 42 41 41 4d
```

```
 Address 0x50: 58 2d 4d 50 43 32 45 2d 33 44 2d 45 51 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff 74 ff ff ff ff ff ff ff ff ff ff ff ff
 CPU            REV 03   711-038484   ZX3665            MPCE PMB 2G
Jedec Code:   0x7fb0            EEPROM Version:    0x01
P/N:          711-038484        S/N:               ZX3665
Assembly ID:  0x0b36            Assembly Version:  01.03
Date:         02-01-2012        Assembly Flags:    0x00
Version:      REV 03
ID: MPCE PMB 2G
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 0b 36 01 03 52 45 56 20 30 33 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 34 38 34 00 00
 Address 0x20: 53 2f 4e 20 5a 58 33 36 36 35 00 00 00 01 02 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff 00 00 00 02 00 00 0c 00 42 5f c0 a4
 MIC 0          REV 05   750-037128   ZR4031            1xCOC12/4xCOC3 CH-CE
Jedec Code:   0x7fb0            EEPROM Version:    0x02
P/N:          750-037128        S/N:               ZR4031
Assembly ID:  0x0a1b            Assembly Version:  01.05
Date:         12-04-2011        Assembly Flags:    0x00
Version:      REV 05            CLEI Code:         PROTOXCLEI
ID: 1xCOC12/4xCOC3 CH-CE        FRU Model Number:  MIC-3D-4CHOC3-1OC12-CE
Board Information Record:
 Address 0x00: 12 01 05 03 05 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0a 1b 01 05 52 45 56 20 30 35 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 31 32 38 00 00
 Address 0x20: 53 2f 4e 20 5a 52 34 30 33 31 00 00 00 04 0c 07
 Address 0x30: db ff ff ff 12 01 05 03 05 ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
 Address 0x50: 49 43 2d 33 44 2d 34 43 48 4f 43 33 2d 31 4f 43
 Address 0x60: 31 32 2d 43 45 00 30 32 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff 98 c0 02 61 bc 7f b0 02 ff 0a 11 01 17
   PIC 0                BUILTIN      BUILTIN           1xCOC12/4xCOC3 CH-CE
 MIC 1          REV 23   750-032479   CADE8614          MIC-3D-8DS3-E3
Jedec Code:   0x7fb0            EEPROM Version:    0x02
P/N:          750-032479        S/N:               CADE8614
Assembly ID:  0x0a11            Assembly Version:  01.23
Date:         07-24-2014        Assembly Flags:    0x00
Version:      REV 23            CLEI Code:         COUIA8DBAA
ID: MIC-3D-8DS3-E3              FRU Model Number:  MIC-3D-8DS3-E3
Board Information Record:
 Address 0x00: 56 01 ff ff 03 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0a 11 01 17 52 45 56 20 32 33 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 33 32 34 37 39 00 00
 Address 0x20: 53 2f 4e 20 43 41 44 45 38 36 31 34 00 18 07 07
 Address 0x30: de ff ff ff 56 01 ff ff 03 ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 01 43 4f 55 49 41 38 44 42 41 41 4d
 Address 0x50: 49 43 2d 33 44 2d 38 44 53 33 2d 45 33 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 41 00 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff 7b c0 03 e5 7c 4f 8a 9e 10 00 00 00 02
   PIC 2                BUILTIN      BUILTIN           MIC-3D-8DS3-E3
 QXM 0          REV 06   711-028408   ZW8299            MPC QXM
Jedec Code:   0x7fb0            EEPROM Version:    0x01
```



```
  P/N:           711-028408        S/N:               ZW8299
  Assembly ID:  0x097a            Assembly Version:  02.06
  Date:         01-19-2012        Assembly Flags:    0x00
  Version:      REV 06
  ID: MPC QXM
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 01 ff 09 7a 02 06 52 45 56 20 30 36 00 00
   Address 0x10: 00 00 00 00 37 31 31 2d 30 32 38 34 30 38 00 00
   Address 0x20: 53 2f 4e 20 5a 57 38 32 39 39 00 00 00 13 01 07
   Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
   Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
   QXM 1           REV 06   711-028408   ZY0609           MPC QXM
  Jedec Code:   0x7fb0            EEPROM Version:    0x01
  P/N:          711-028408        S/N:               ZY0609
  Assembly ID:  0x097a            Assembly Version:  02.06
  Date:         01-19-2012        Assembly Flags:    0x00
  Version:      REV 06
  ID: MPC QXM
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 01 ff 09 7a 02 06 52 45 56 20 30 36 00 00
   Address 0x10: 00 00 00 00 37 31 31 2d 30 32 38 34 30 38 00 00
   Address 0x20: 53 2f 4e 20 5a 59 30 36 30 39 00 00 00 13 01 07
   Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
   Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
ADC 7             REV 17   750-043596   ABCA0990          Adapter Card
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          750-043596        S/N:               ABCA0990
  Assembly ID:  0x0b3d            Assembly Version:  01.17
  Date:         03-07-2013        Assembly Flags:    0x00
  Version:      REV 17            CLEI Code:         IPUCBA8CAA
  ID: Adapter Card                FRU Model Number:  MX2000-LC-ADAPTER
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  I2C Hex Data:
   Address 0x00: 7f b0 02 ff 0b 3d 01 11 52 45 56 20 31 37 00 00
   Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
   Address 0x20: 53 2f 4e 20 41 42 43 41 30 39 39 30 00 07 03 07
   Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
   Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
   Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
   Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
   Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00 00
Fan Tray 0        REV 01   760-052467   ACAY6190          172mm FanTray - 6 Fans
  Jedec Code:   0x7fb0            EEPROM Version:    0x02
  P/N:          760-052467        S/N:               ACAY6190
  Assembly ID:  0x0b96            Assembly Version:  02.10
  Date:         09-18-2015        Assembly Flags:    0x00
  Version:      REV 01            CLEI Code:         IPUCBENCAA
  ID: 172mm FanTray - 6 Fans      FRU Model Number:  MX2000-FANTRAY-S
  Board Information Record:
   Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```



```
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 96 02 0a 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 35 32 34 36 37 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 59 36 31 39 30 00 12 09 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 43 42 45 4e 43 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 46 41 4e 54 52 41 59 2d 53 00
  Address 0x60: 00 00 00 00 00 00 31 ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 1a ff ff ff ff ff ff ff ff ff ff ff ff
Fan Tray 1       REV 01   760-052467   ACAY6414         172mm FanTray - 6 Fans
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          760-052467        S/N:               ACAY6414
 Assembly ID:  0x0b96            Assembly Version:  02.10
 Date:         10-28-2015        Assembly Flags:    0x00
 Version:      REV 01            CLEI Code:         IPUCBENCAA
 ID: 172mm FanTray - 6 Fans      FRU Model Number:  MX2000-FANTRAY-S
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 96 02 0a 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 35 32 34 36 37 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 59 36 34 31 34 00 1c 0a 07
  Address 0x30: df ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 43 42 45 4e 43 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 46 41 4e 54 52 41 59 2d 53 00
  Address 0x60: 00 00 00 00 00 00 31 ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 1a ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware models (MX2008 Router)

```
user@host>show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 47   750-044636   ABAD1739          CHAS-BP-MX2010-S
PMP              REV 01   711-051406   ACVD0738
FPM Board        REV 02   760-068193   ABDG7408          PROTO-ASSEMBLY
PSM 1            REV 06   740-050037   1EDB61200R8       MX2000-PSM-DC-S
PSM 2            REV 06   740-050037   1EDB61200WA       MX2000-PSM-DC-S
PSM 3            REV 06   740-050037   1EDB61200NY       MX2000-PSM-DC-S
PSM 4            REV 06   740-050037   1EDB61200N2       MX2000-PSM-DC-S
PSM 5            REV 06   740-050037   1EDB61200RN       MX2000-PSM-DC-S
PSM 6            REV 06   740-050037   1EDB61200RF       MX2000-PSM-DC-S
PSM 7            REV 06   740-050037   1EDB61200R7       MX2000-PSM-DC-S
PDM 0            REV 01   740-060189   1EFF5250143       MX2K-PDM-OP-DC-S
PDM 1            REV 01   740-060189   1EFF5250074       MX2K-PDM-OP-DC-S
CB 0             REV 01   750-067373   ABDJ0047          PROTO-ASSEMBLY
CB 1             REV 03   750-067373   ABDH3016          PROTO-ASSEMBLY
SFB 0            REV 08   750-067371   ABDK7180          PROTO-ASSEMBLY
SFB 1            REV 08   750-067371   ABDK7024          PROTO-ASSEMBLY
SFB 2            REV 08   750-067371   ABDK7188          PROTO-ASSEMBLY
SFB 3            REV 08   750-067371   ABDK7143          PROTO-ASSEMBLY
SFB 4            REV 08   750-067371   ABDK7030          PROTO-ASSEMBLY
SFB 5            REV 08   750-067371   ABDK7146          PROTO-ASSEMBLY
SFB 6            REV 08   750-067371   ABDK7203          PROTO-ASSEMBLY
SFB 7            REV 08   750-067371   ABDK7238          PROTO-ASSEMBLY
FPC 0            REV 36   750-044130   ABCS8607          PROTO-ASSEMBLY
  MIC 0          REV 21   750-050008   ABCT5920          MIC6-100G-CXP
FPC 1            REV 22   750-063414   CAFJ3026          MX2K-MPC9E
FPC 7            REV 08   750-038492   ZX4080            MX-MPC2E-3D-EQ
  MIC 0          REV 05   750-037128   ZR4031            MIC-3D-4CHOC3-1OC12-CE
  MIC 1          REV 23   750-032479   CADE8614          MIC-3D-8DS3-E3
```

```
ADC 7            REV 17   750-043596   ABCA0990          MX2000-LC-ADAPTER
Fan Tray 0       REV 01   760-052467   ACAY6190          MX2000-FANTRAY-S
Fan Tray 1       REV 01   760-052467   ACAY6414          MX2000-FANTRAY-S
```

show chassis hardware clei-models (MX2008 Router)

```
user@host>show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 47   750-044636   IPMU810ARB        CHAS-BP-MX2010-S
PMP              REV 01   711-051406
FPM Board        REV 02   760-068193   PROTOXCLEI        PROTO-ASSEMBLY
PSM 1            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PSM 2            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PSM 3            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PSM 4            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PSM 5            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PSM 6            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PSM 7            REV 06   740-050037   IPUPAPDKAA        MX2000-PSM-DC-S
PDM 0            REV 01   740-060189   IPUPAN1KAA        MX2K-PDM-OP-DC-S
PDM 1            REV 01   740-060189   IPUPAN1KAA        MX2K-PDM-OP-DC-S
CB 0             REV 01   750-067373   PROTOXCLEI        PROTO-ASSEMBLY
CB 1             REV 03   750-067373   PROTOXCLEI        PROTO-ASSEMBLY
SFB 0            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 1            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 2            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 3            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 4            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 5            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 6            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
SFB 7            REV 08   750-067371   PROTOXCLEI        PROTO-ASSEMBLY
FPC 0            REV 36   750-044130   PROTOXCLEI        PROTO-ASSEMBLY
  MIC 0          REV 21   750-050008   IP9IATYDAA        MIC6-100G-CXP
FPC 1            REV 22   750-063414   IPUCBMUCAA        MX2K-MPC9E
FPC 7            REV 08   750-038492   COUIBA5BAA        MX-MPC2E-3D-EQ
  MIC 0          REV 05   750-037128   PROTOXCLEI        MIC-3D-4CHOC3-1OC12-CE
  MIC 1          REV 23   750-032479   COUIA8DBAA        MIC-3D-8DS3-E3
ADC 7            REV 17   750-043596   IPUCBA8CAA        MX2000-LC-ADAPTER
Fan Tray 0       REV 01   760-052467   IPUCBENCAA        MX2000-FANTRAY-S
Fan Tray 1       REV 01   760-052467   IPUCBENCAA        MX2000-FANTRAY-S
```

show chassis hardware (MX10003 Router)

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BLANK             JNP10003 [MX10003]
Midplane         REV 01   750-066883   CAGM0759          Midplane 2
Routing Engine 0          BUILTIN      BUILTIN           Routing Engine
Routing Engine 1          BUILTIN      BUILTIN           Routing Engine
CB 0             REV 07   750-067071   CAGX4354          SPM
  Mezz           REV 10   711-066896   CAHS7200          SPM Mezz Board
CB 1             REV 07   750-067071   CAGX4363          SPM
  Mezz           REV 10   711-066896   CAHS7193          SPM Mezz Board
FPC 0            REV 05   750-066879   CAGV0273          LC2103
  CPU                     BUILTIN      BUILTIN           SMPC PMB
  PIC 0
  PIC 1
FPC 1            REV 05   750-066879   CAGV0278          LC2103
```



```
  CPU                        BUILTIN      BUILTIN           SMPC PMB
  PIC 0                      BUILTIN      BUILTIN           6xQSFPP
  PIC 1
PEM 0            REV 01     740-066937   1HS16320003       JNP-PWR1600-AC
PEM 1            REV 01     740-066937   1HS16320002       JNP-PWR1600-AC
Fan Tray 0       REV 02     760-069329   CAGS7731          JNP FAN 3RU
Fan Tray 1       REV 02     760-069329   CAGS7776          JNP FAN 3RU
Fan Tray 2       REV 02     760-069329   CAGS7659          JNP FAN 3RU
Fan Tray 3       REV 02     760-069329   CAGS7669          JNP FAN 3RU
```

show chassis hardware (MX204 Router)

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BB768             JNP204 [MX204]
Routing Engine 0          BUILTIN      BUILTIN           RE-S-2X00x6
CB 0             REV 11   750-069579   CAJD3113          JNP204 [MX204]
FPC 0                     BUILTIN      BUILTIN           MPC
  PIC 0                   BUILTIN      BUILTIN           4XQSFP28 PIC
    Xcvr 0       REV 01   740-061405   1ACQ110409R       QSFP-100GBASE-SR4
    Xcvr 1       REV 01   740-054053   QF027546          QSFP+-4X10G-SR
    Xcvr 2       REV 01   740-058732   1AMQA142092       QSFP-100GBASE-LR4
    Xcvr 3       REV 01   740-058732   1AMQA14203J       QSFP-100GBASE-LR4
  PIC 1                   BUILTIN      BUILTIN           8XSFPP PIC
PEM 1            REV 04   740-043886   1GA46361256       JPSU-650W-DC-AFO
Fan Tray 0                                               Fan Tray, Front to Back
 Airflow - AFO
Fan Tray 1                                               Fan Tray, Front to Back
 Airflow - AFO
Fan Tray 2                                               Fan Tray, Front to Back
 Airflow - AFO
```

show chassis hardware (vMX running in lite mode)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VM54599D128A      VMX
Midplane
Routing Engine 0                                         RE-VMX
CB 0                                                     VMX SCB
CB 1                                                     VMX SCB
FPC 0                                                    Virtual FPC
  CPU            Rev. 1.0 RIOT-LITE    BUILTIN
  MIC 0                                                  Virtual
    PIC 0                 BUILTIN      BUILTIN           Virtual
```

show chassis hardware (vMX running in performance mode)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VM54599D128A      VMX
Midplane
Routing Engine 0                                         RE-VMX
CB 0                                                     VMX SCB
CB 1                                                     VMX SCB
FPC 0                                                    Virtual FPC
```

```
  CPU            Rev. 1.0 RIOT-PERF    BUILTIN
  MIC 0                                                 Virtual
    PIC 0                 BUILTIN      BUILTIN          Virtual
```

show chassis hardware (T320 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                19093             T320
Midplane         REV 04   710-004339   BC1436            T320 Backplane
FPM GBUS         REV 03   710-004461   BC1407            T320 FPM Board
FPM Display      REV 04   710-002897   BE0763            FPM Display
CIP              REV 05   710-002895   BB2311            T Series CIP
PEM 0            Rev 01   740-004359   NB12546           Power Entry Module
SCG 0            REV 06   710-004455   AY4522            T320 Sonet
Clock Gen.
Routing Engine 0                                         unknown
CB 0             REV 13   710-002728   BC1577            T Series
Control Board
CB 1             REV 13   710-002728   BC1595            T Series
Control Board
FPC 1            REV 09   710-007531   HS1572            FPC Type 2
   CPU            REV 15   710-001726   HR8763            FPC CPU
   PIC 0          REV 01   750-010618   CB5579            4x G/E SFP,
1000 BASE
     SFP 0        REV 01   740-007326   P5809Z1           SFP-SX
     SFP 1        REV 01   740-007326   P4Q10XU           SFP-SX
     SFP 2                 NON-JNPR     RA45020031        SFP-SX
     SFP 3                 NON-JNPR     RA45020032        SFP-SX
   PIC 1          REV 01   750-010618   CD9587            4x G/E SFP,
1000 BASE
     SFP 0                 NON-JNPR     P5A08QZ           SFP-T
     SFP 1        REV 01   740-007326   P4Q133K           SFP-SX
     SFP 2        REV 01   740-007326   P5809YY           SFP-SX
     SFP 3        REV 01   740-007327   4C81704           SFP-LX
   MMB 1          REV 03   710-005555   HR9401            MMB-288mbit
   PPB 0          REV 04   710-003758   HR2886            PPB Type 2
FPC 2            REV 07   710-005860   HP2392            FPC Type 1
   CPU            REV 14   710-001726   HP7797            FPC CPU
   PIC 0          REV 02   750-007643   HM0853            1x G/E QPP,
1000 BASE
     SFP 0        REV 01   740-007326   P11E9JJ           SFP-SX
   MMB 1          REV 02   710-005555   HN2379            MMB-288mbit
   PPB 0          REV 04   710-003758   HP8092            PPB Type 2
FPC 3            REV 07   710-005860   HP2393            FPC Type 1
   CPU            REV 14   710-001726   HP0968            FPC CPU
   PIC 0          REV 01   750-010240   CB5363            1x G/E SFP,
1000 BASE
     SFP 0        REV 01   740-007326   P4R0PNH           SFP-SX
   PIC 1          REV 03   750-003034   HD2832            4x OC-3 SONET,
SMIR
   MMB 1          REV 02   710-005555   HN6307            MMB-288mbit
   PPB 0          REV 04   710-003758   HP5051            PPB Type 2
FPC 4            REV 01   710-010845   JD3872            FPC Type 4
   CPU            REV 02   710-011481   JB6042            FPC CPU
 5            REV 01   710-005802   BC1566            FPC Type 2
   CPU            REV 09   710-001726   AY4922            FPC CPU
   PIC 0          REV 02   750-008155   BE2114            2x G/E QPP,
1000 BASE
     SFP 0        REV 01   740-007326   P4R0PMQ           SFP-SX
```



```
     SFP 1         REV 01    740-007326    P4R0PN9          SFP-SX
   PIC 1           REV 01    750-008155    BE2116           2x G/E QPP,
1000 BASE
     SFP 0         REV 01    740-007326    P4R0PNZ          SFP-SX
     SFP 1                   NON-JNPR      2908             SFP-T
   MMB 1           REV 01    710-005555    AZ2246           MMB-288mbit
   PPB 0           REV 03    710-003758    AY4839           PPB Type 2
FPC 7             REV 01    710-005803    AZ2123           FPC Type 3
...
```

show chassis hardware (T640 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                19182             T640
Midplane         REV 04   710-002726   AX5608            T640 Backplane
FPM GBUS         REV 02   710-002901   HE3064            T640 FPM Board
FPM Display      REV 02   710-002897   HE7864            FPM Display
CIP              REV 05   710-002895   HA5024            T Series CIP
PEM 0            Rev 02   740-029522   VH26235           AC PEM 10kW US
PEM 1            Rev 02   740-029522   VH26230           AC PEM 10kW US
SCG 0            REV 03   710-003423   HA4508            T640 Sonet Clock Gen.
Routing Engine 0 REV 02   740-005022   210865700483      RE-3.0 (RE-600)
CB 0             REV 01   710-002728   HD3044            T Series Control Board
FPC 2            REV 04   710-001721   HD5572            FPC Type 3
  CPU            REV 06   710-001726   HA4712            FPC CPU
  PIC 1          REV 03   750-009567   HV2331            1x 10GE(LAN),XENPAK
    SFP 0        REV 01   740-009898   USC202R103        XENPAK-SR
  PIC 2          REV 03   750-009567   HV2332            1x 10GE(LAN),XENPAK
    SFP 0        REV 01   740-011268   USC202R112        XENPAK-ZR
  PIC 3          REV 03   750-009567   HX4416            1x 10GE(LAN),XENPAK
    SFP 0        REV 01   740-012056   434TC004          XENPAK-CX4
  PIC 4          REV 03   750-009567   HX4420            1x 10GE(LAN),XENPAK
    SFP 0        REV 01   740-012058   434TC124          XENPAK-LX4
FPC 5            REV 01   710-013553   JE4839            E2-FPC Type 1
  CPU            REV 01   710-013569   JW9163            FPC CPU
  PIC 0          REV 01   750-009567   HX4419            1x 10GE(LAN),XENPAK
    SFP 0        REV 01   740-009898   USC202RT05        XENPAK-LR
  PIC 1          REV 03   750-009567   HN7426            1x 10GE(LAN),XENPAK
    SFP 0        REV 01   740-009550   03L90051          XENPAK-ER
  PIC 2          REV 03   750-009467   HT7423            1x 10GE(LAN),XENPAK
    SFP 0                 NON-JNPR                       UNKNOWON
  PIC 3          REV 04   750-005100   AY4850            1x 10GE(LAN),DWDM
FPC 4            REV 01   710-010845   JD3872            FPC Type 4
  CPU            REV 02   710-011481   JB6042            FPC CPU
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Fan Tray
```

show chassis hardware models (T640 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 04   710-002726                     CHAS-BP-T640-S
FPM Display      REV 02   710-002897                     CRAFT-T640-S
CIP              REV 05   710-002895                     CIP-L-T640-S
PEM 0            Rev 01   740-002595                     PWR-T-DC-S
SCG 0            REV 04   710-003423                     SCG-T-S
```



```
SCG 1            REV 04   710-003423                     SCG-T-S
Routing Engine 0 REV 01   740-005022                     RE-600-2048-S
Routing Engine 1 REV 07   740-005022                     RE-600-2048-S
CB 0             REV 06   710-002726                     CHAS-BP-T640-S
CB 1             REV 06   710-002728                     CB-L-T-S
FPC 5            REV 05   710-007527                     T640-FPC2
  PIC 0          REV 05   750-002510                     PB-2GE-SX
  PIC 1          REV 05   750-001901                     PB-4OC12-SON-SMIR
FPC 6            REV 03   710-001721                     T640-FPC3
  PIC 1          REV 01   750-009553                     PC-4OC48-SON-SFP
SIB 4            REV 02   750-005486                     SIB-I-T640-S
Fan Tray 0                                               FANTRAY-T-S
Fan Tray 1                                               FANTRAY-T-S
Fan Tray 2                                               FAN-REAR-TX-T640-S
```

show chassis hardware extensive (T640 Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                                  T640
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:        ...........       S/N:             ...........
 Assembly ID:  0x0507            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 Version:    ...........
 ID: Gibson LCC Chassis
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 04   710-002726   AX5633
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          710-002726.       S/N:               AX5633.
 Assembly ID:  0x0127            Assembly Version:  01.04
 Date:         06-27-2001        Assembly Flags:    0x00
 Version:      REV 04.....
 ID: Gibson Backplane
 Board Information Record:
  Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00
  Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07
  Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff
  Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM GBUS         REV 02   710-002901   HE3245
 ...
FPM Display      REV 02   710-002897   HA4873
 ...
CIP              REV 05   710-002895   HA4729
...
PEM 1            RevX02   740-002595   MD21815           Power Entry Module
 ...
SCG 0            REV 04   710-003423   HF6023
 ...
SCG 1            REV 04   710-003423   HF6061
```

```
 ...
Routing Engine 0 REV 01   740-005022   210865700292      RE-3.0
...
CB 0             REV 06   710-002728   HE3614
...
FPC 1            REV 01   710-002385   HE3009            FPC Type 1
...              REV 06   710-001726   HC0010
```

show chassis hardware (T4000 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1172F25AHA      T4000
Midplane         REV 01   710-027486   RC8355            T-series Backplane
FPM GBUS         REV 13   710-002901   BBAE0927          T640 FPM Board
FPM Display      REV 01   710-021387   EF6764            T1600 FPM Display
CIP              REV 06   710-002895   BBAD9210          T-series CIP
PEM 0            REV 01   740-036442   VA00016           Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAD7248          T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAE3874          T640 Sonet Clock Gen.
Routing Engine 0 REV 05   740-026941   P737F-002248      RE-DUO-1800
Routing Engine 1 REV 06   740-026941   P737F-002653      RE-DUO-1800
CB 0             REV 09   710-022597   ED0295            LCC Control Board
CB 1             REV 09   710-022597   EA6050            LCC Control Board
FPC 0            REV 26   750-032819   EK1173            FPC Type 5-3D
  CPU            REV 12   711-030686   EJ8584            SNG PMB
  PIC 0          REV 07   750-034624   EF6837            12x10GE (LAN/WAN) SFPP
    Xcvr 0       REV 01   740-031980   123363A01145      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   123363A01147      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJJ01P3           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B10M03256         SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   AJJ01M2           SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   123363A01137      SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   AJJ01PN           SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   AJJ01NW           SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   123363A01139      SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   AJJ01KE           SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   123363A01336      SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   B10M01325         SFP+-10G-SR
  PIC 1          REV 07   750-034624   EF6800            12x10GE (LAN/WAN) SFPP
    Xcvr 0       REV 01   740-031980   AJJ01SA           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJJ01QZ           SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJH0217           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AJJ01TE           SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   AJJ01KV           SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   AJJ01MU           SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   AJJ01R0           SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   AJJ01TC           SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   AJJ0364           SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   AJD0GV3           SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   B10M03343         SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   AJJ01QJ           SFP+-10G-SR
  LMB 0          REV 05   711-034381   EJ8490            Type-0 LMB
  LMB 1          REV 04   711-035774   EJ8517            Type-1 LMB
  LMB 2          REV 05   711-034381   EJ8489            Type-0 LMB
FPC 3            REV 07   750-032819   EG3637            FPC Type 5-3D
  CPU            REV 09   711-030686   EG0150            SNG PMB
  PIC 0          REV 08   750-035293   EF3657            1x100GE
    Xcvr 0       REV 01   740-032210   C22CQNJ           CFP-100G-LR4
  PIC 1          REV 10   750-034624   BBAN4098          12x10GE (LAN/WAN) SFPP
```



```
    Xcvr 0       REV 01   740-031980   B11J04902        SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11J04891        SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJJ01MX          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11J04183        SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   B11J04894        SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   B11J04184        SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   B11J04897        SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   B11J04899        SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   AJJ01TV          SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   B11J04057        SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   AJJ01M4          SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   B11J04905        SFP+-10G-SR
  LMB 0          REV 04   711-034381   EG1524           Type-0 LMB
  LMB 1          REV 03   711-035774   EG0345           Type-1 LMB
  LMB 2          REV 04   711-034381   EG1522           Type-0 LMB
FPC 5            REV 03   710-033871   BBAJ0768         FPC Type 4-ES
  CPU            REV 11   710-016744   BBAH9342         ST-PMB2
  PIC 0          REV 09   750-029262   EE6789           100GE
  PIC 1          REV 03   750-034781   EE6655           100GE CFP
    Xcvr 0       REV 01   740-032210   J11A22334        CFP-100G-LR4
  BRIDGE 0       REV 03   711-029995   EE6572           100GE Bridge Board
  MMB 0          REV 07   710-025563   BBAJ4657         ST-MMB2
  MMB 1          REV 07   710-025563   BBAJ3073         ST-MMB2
FPC 6            REV 05   750-010153   EF4936           FPC Type 5-3D
  CPU            REV 06   711-030686   EF4189           SNG PMB
  PIC 0          REV 10   750-034624   BBAN4109         12x10GE (LAN/WAN) SFPP
    Xcvr 0       REV 01   740-031980   B11J04895        SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11J04898        SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   B11J04021        SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B11J04903        SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   B11J04311        SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   B11J04059        SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   B11J04016        SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   B11J04017        SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   B11J04887        SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   B11J04297        SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   B11J04893        SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   B11J04022        SFP+-10G-SR
  PIC 1          REV 02   750-034624   EE3711           12x10GE (LAN/WAN) SFPP
    Xcvr 0       REV 01   740-031980   AJH033X          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJJ01N0          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJJ01SV          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AJJ032L          SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   B10M01593        SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   AJD0FF1          SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   AJJ01NU          SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   123363A01305     SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   B10M00361        SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   AJJ01M7          SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   AJJ032X          SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   AJJ01PG          SFP+-10G-SR
  LMB 0          REV 04   711-034381   EF3838           Type-0 LMB
  LMB 1          REV 03   711-035774   EF3821           Type-1 LMB
  LMB 2          REV 04   711-034381   EF3834           Type-0 LMB
SPMB 0           REV 05   710-023321   ED1990           LCC Switch CPU
SPMB 1           REV 05   710-023321   EA2768           LCC Switch CPU
SIB 0            REV 02   711-036340   EF8802           SIB-HC-3D
SIB 1            REV 07   711-036340   EG2286           SIB-HC-3D
SIB 2            REV 07   711-036340   EG2252           SIB-HC-3D
SIB 3            REV 02   711-036340   EF1358           SIB-HC-3D
SIB 4            REV 02   711-036340   EF8806           SIB-HC-3D
```



```
Fan Tray 0                                                  Front Top Fan Tray
Fan Tray 1                                                  Front Bottom Fan Tray
-- Rev 2
Fan Tray 2                                                  Rear Fan Tray -- Rev 3
```

show chassis hardware (T4000 Router with 16-GB Line Card Chassis (LCC) Routing Engine)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11BDF2CAHA      T1600
Midplane         REV 01   710-027486   ACAJ0774          T640 Backplane
FPM GBUS         REV 13   710-002901   BBAL6812          T640 FPM Board
FPM Display      REV 04   710-021387   BBAP2679          T1600 FPM Display
CIP              REV 06   710-002895   BBAP4758          T-series CIP
PEM 0            Rev 03   740-026384   XF86421           Power Entry Module 3x80
PEM 1            Rev 03   740-026384   XF86429           Power Entry Module 3x80
SCG 0            REV 18   710-003423   BBAP1896          T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAN8659          T640 Sonet Clock Gen.
Routing Engine 0 REV 01   740-042243   737F-002238       RE-DUO-1800-16G
Routing Engine 1 REV 01   740-042243   737F-002403       RE-DUO-1800-16G
CB 1             REV 11   710-022597   EK4526            LCC Control Board
CB 1             REV 11   710-022597   EK4527            LCC Control Board
FPC 0            REV 05   710-033871   EK5644            FPC Type 4-ES
  CPU            REV 11   710-016744   EK3428            ST-PMB2
  PIC 0          REV 20   750-017405   EJ3041            4x 10GE (LAN/WAN) XFP
  PIC 1          REV 17   750-026962   EH7536            10x10GE(LAN/WAN) SFPP
  MMB 0          REV 07   710-025563   EK6039            ST-MMB2
  MMB 1          REV 07   710-025563   EK6086            ST-MMB2
FPC 1            REV 05   710-033871   EK6583            FPC Type 4-ES
  CPU            REV 11   710-016744   EK3401            ST-PMB2
  PIC 0          REV 17   750-026962   EJ8948            10x10GE(LAN/WAN) SFPP
  MMB 0          REV 07   710-025563   EK6202            ST-MMB2
  MMB 1          REV 07   710-025563   EK6112            ST-MMB2
SPMB 1           REV 05   710-023321   EK4900            LCC Switch CPU
SIB 0            REV 11   710-013074   EK5958            SIB-I8-SF
SIB 1            REV 11   710-013074   EK4606            SIB-I8-SF
SIB 2            REV 11   710-013074   EK5971            SIB-I8-SF
SIB 3            REV 11   710-013074   EK4609            SIB-I8-SF
SIB 4            REV 11   710-013074   EK4602            SIB-I8-SF
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Fan Tray -- Rev 2
```

show chassis hardware (T4000 Router with LSR FPC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1173A24AHA      T4000
FPC 3            REV      750-048373   AN7797            FPC Type 5-LSR
  CPU            REV 10   711-030686   AN6649            SNG PMB
  PIC 0          REV 07   750-034624   EF6830            12x10GE (LAN/WAN) SFPP
```

show chassis hardware clei-models (T4000 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   710-027486   IPMJ700DRD        CHAS-BP-T1600-S
```

```
FPM Display      REV 01   710-021387                    CRAFT-T1600-S
CIP              REV 06   710-002895                    CIP-L-T640-S
PEM 0            REV 01   740-036442   IPUPAG6KAA       PWR-T-6-60-DC
SCG 0            REV 18   710-003423                    SCG-T-S
SCG 1            REV 18   710-003423                    SCG-T-S
Routing Engine 0 REV 05   740-026941                    RE-DUO-C1800-8G-S
Routing Engine 1 REV 06   740-026941                    RE-DUO-C1800-8G-S
CB 0             REV 09   710-022597                    CB-LCC-S
CB 1             REV 09   710-022597                    CB-LCC-S
FPC 3
  PIC 0          REV 08   750-035293   XXXXXXXXBB       PF-1CGE-CFP
  PIC 1          REV 10   750-034624   XXXXXXXXCC       PF-12XGE-SFPP
FPC 5            REV 03   710-033871   IPUCAMBCTD       T1600-FPC4-ES
  PIC 1          REV 03   750-034781   IPUIBKLMAA       PD-1CE-CFP-FPC4
FPC 6
  PIC 0          REV 10   750-034624   XXXXXXXXCC       PF-12XGE-SFPP
Fan Tray 0                                              FANTRAY-T-S
Fan Tray 1                                              FANTRAY-T4000-S
Fan Tray 2                                              FANTRAY-TXP-R-S
```

show chassis hardware detail (T4000 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN1172F25AHA     T4000
Midplane         REV 01   710-027486   RC8355           T-series Backplane
FPM GBUS         REV 13   710-002901   BBAE0927         T640 FPM Board
FPM Display      REV 01   710-021387   EF6764           T1600 FPM Display
CIP              REV 06   710-002895   BBAD9210         T-series CIP
PEM 0            REV 01   740-036442   VA00016          Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAD7248         T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAE3874         T640 Sonet Clock Gen.
Routing Engine 0 REV 05   740-026941   P737F-002248     RE-DUO-1800
  ad0     3823 MB  SMART CF            2009121602A661576157 Compact Flash
  ad1    59690 MB  STEC MACH-8 SSD     STM000103FDB     Disk 1
Routing Engine 1 REV 06   740-026941   P737F-002653     RE-DUO-1800
  ad0     3823 MB  SMART CF            201011150153F52CF52C Compact Flash
  ad1    62720 MB  SMART Lite SATA Drive 2010110900150A880A88 Disk 1
CB 0             REV 09   710-022597   ED0295           LCC Control Board
CB 1             REV 09   710-022597   EA6050           LCC Control Board
FPC 0            REV 26   750-032819   EK1173           FPC Type 5-3D
  CPU            REV 12   711-030686   EJ8584           SNG PMB
  PIC 0          REV 07   750-034624   EF6837           12x10GE (LAN/WAN) SFPP
    Xcvr 0       REV 01   740-031980   123363A01145     SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   123363A01147     SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJJ01P3          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   B10M03256        SFP+-10G-SR
    Xcvr 4       REV 01   740-031980   AJJ01M2          SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   123363A01137     SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   AJJ01PN          SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   AJJ01NW          SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   123363A01139     SFP+-10G-SR
    Xcvr 9       REV 01   740-031980   AJJ01KE          SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   123363A01336     SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   B10M01325        SFP+-10G-SR
  PIC 1          REV 07   750-034624   EF6800           12x10GE (LAN/WAN) SFPP
    Xcvr 0       REV 01   740-031980   AJJ01SA          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AJJ01QZ          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AJH0217          SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AJJ01TE          SFP+-10G-SR
```



```
      Xcvr 4        REV 01   740-031980   AJJ01KV          SFP+-10G-SR
      Xcvr 5        REV 01   740-031980   AJJ01MU          SFP+-10G-SR
      Xcvr 6        REV 01   740-031980   AJJ01R0          SFP+-10G-SR
      Xcvr 7        REV 01   740-031980   AJJ01TC          SFP+-10G-SR
      Xcvr 8        REV 01   740-031980   AJJ0364          SFP+-10G-SR
      Xcvr 9        REV 01   740-031980   AJD0GV3          SFP+-10G-SR
      Xcvr 10       REV 01   740-031980   B10M03343        SFP+-10G-SR
      Xcvr 11       REV 01   740-031980   AJJ01QJ          SFP+-10G-SR
    LMB 0           REV 05   711-034381   EJ8490           Type-0 LMB
    LMB 1           REV 04   711-035774   EJ8517           Type-1 LMB
    LMB 2           REV 05   711-034381   EJ8489           Type-0 LMB
  FPC 3             REV 07   750-032819   EG3637           FPC Type 5-3D
    CPU             REV 09   711-030686   EG0150           SNG PMB
    PIC 0           REV 08   750-035293   EF3657           1x100GE
      Xcvr 0        REV 01   740-032210   C22CQNJ          CFP-100G-LR4
    PIC 1           REV 10   750-034624   BBAN4098         12x10GE (LAN/WAN) SFPP
      Xcvr 0        REV 01   740-031980   B11J04902        SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   B11J04891        SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   AJJ01MX          SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   B11J04183        SFP+-10G-SR
      Xcvr 4        REV 01   740-031980   B11J04894        SFP+-10G-SR
      Xcvr 5        REV 01   740-031980   B11J04184        SFP+-10G-SR
      Xcvr 6        REV 01   740-031980   B11J04897        SFP+-10G-SR
      Xcvr 7        REV 01   740-031980   B11J04899        SFP+-10G-SR
      Xcvr 8        REV 01   740-031980   AJJ01TV          SFP+-10G-SR
      Xcvr 9        REV 01   740-031980   B11J04057        SFP+-10G-SR
      Xcvr 10       REV 01   740-031980   AJJ01M4          SFP+-10G-SR
      Xcvr 11       REV 01   740-031980   B11J04905        SFP+-10G-SR
    LMB 0           REV 04   711-034381   EG1524           Type-0 LMB
    LMB 1           REV 03   711-035774   EG0345           Type-1 LMB
    LMB 2           REV 04   711-034381   EG1522           Type-0 LMB
  FPC 5             REV 03   710-033871   BBAJ0768         FPC Type 4-ES
    CPU             REV 11   710-016744   BBAH9342         ST-PMB2
    PIC 0           REV 09   750-029262   EE6789           100GE
    PIC 1           REV 03   750-034781   EE6655           100GE CFP
      Xcvr 0        REV 01   740-032210   J11A22334        CFP-100G-LR4
    BRIDGE 0        REV 03   711-029995   EE6572           100GE Bridge Board
    MMB 0           REV 07   710-025563   BBAJ4657         ST-MMB2
    MMB 1           REV 07   710-025563   BBAJ3073         ST-MMB2
  FPC 6             REV 05   750-010153   EF4936           FPC Type 5-3D
    CPU             REV 06   711-030686   EF4189           SNG PMB
    PIC 0           REV 10   750-034624   BBAN4109         12x10GE (LAN/WAN) SFPP
      Xcvr 0        REV 01   740-031980   B11J04895        SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   B11J04898        SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   B11J04021        SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   B11J04903        SFP+-10G-SR
      Xcvr 4        REV 01   740-031980   B11J04311        SFP+-10G-SR
      Xcvr 5        REV 01   740-031980   B11J04059        SFP+-10G-SR
      Xcvr 6        REV 01   740-031980   B11J04016        SFP+-10G-SR
      Xcvr 7        REV 01   740-031980   B11J04017        SFP+-10G-SR
      Xcvr 8        REV 01   740-031980   B11J04887        SFP+-10G-SR
      Xcvr 9        REV 01   740-031980   B11J04297        SFP+-10G-SR
      Xcvr 10       REV 01   740-031980   B11J04893        SFP+-10G-SR
      Xcvr 11       REV 01   740-031980   B11J04022        SFP+-10G-SR
    PIC 1           REV 02   750-034624   EE3711           12x10GE (LAN/WAN) SFPP
      Xcvr 0        REV 01   740-031980   AJH033X          SFP+-10G-SR
      Xcvr 1        REV 01   740-031980   AJJ01N0          SFP+-10G-SR
      Xcvr 2        REV 01   740-031980   AJJ01SV          SFP+-10G-SR
      Xcvr 3        REV 01   740-031980   AJJ032L          SFP+-10G-SR
      Xcvr 4        REV 01   740-031980   B10M01593        SFP+-10G-SR
      Xcvr 5        REV 01   740-031980   AJD0FF1          SFP+-10G-SR
```



```
    Xcvr 6        REV 01   740-031980   AJJ01NU          SFP+-10G-SR
    Xcvr 7        REV 01   740-031980   123363A01305     SFP+-10G-SR
    Xcvr 8        REV 01   740-031980   B10M00361        SFP+-10G-SR
    Xcvr 9        REV 01   740-031980   AJJ01M7          SFP+-10G-SR
    Xcvr 10       REV 01   740-031980   AJJ032X          SFP+-10G-SR
    Xcvr 11       REV 01   740-031980   AJJ01PG          SFP+-10G-SR
  LMB 0           REV 04   711-034381   EF3838           Type-0 LMB
  LMB 1           REV 03   711-035774   EF3821           Type-1 LMB
  LMB 2           REV 04   711-034381   EF3834           Type-0 LMB
SPMB 0            REV 05   710-023321   ED1990           LCC Switch CPU
SPMB 1            REV 05   710-023321   EA2768           LCC Switch CPU
SIB 0             REV 02   711-036340   EF8802           SIB-HC-3D
SIB 1             REV 07   711-036340   EG2286           SIB-HC-3D
SIB 2             REV 07   711-036340   EG2252           SIB-HC-3D
SIB 3             REV 02   711-036340   EF1358           SIB-HC-3D
SIB 4             REV 02   711-036340   EF8806           SIB-HC-3D
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
-- Rev 2
Fan Tray 2                                               Rear Fan Tray -- Rev 3
```

show chassis hardware models (T4000 Router)

user@host> **show chassis hardware models**

```
Hardware inventory:
Item              Version  Part number  Serial number    FRU model number
Midplane          REV 01   710-027486   RC8355           CHAS-BP-T1600-S
FPM Display       REV 01   710-021387   EF6764           CRAFT-T1600-S
CIP               REV 06   710-002895   BBAD9210         CIP-L-T640-S
PEM 0             REV 01   740-036442   VA00016          PWR-T-6-60-DC
SCG 0             REV 18   710-003423   BBAD7248         SCG-T-S
SCG 1             REV 18   710-003423   BBAE3874         SCG-T-S
Routing Engine 0 REV 05    740-026941   P737F-002248     RE-DUO-C1800-8G-S
Routing Engine 1 REV 06    740-026941   P737F-002653     RE-DUO-C1800-8G-S
CB 0              REV 09   710-022597   ED0295           CB-LCC-S
CB 1              REV 09   710-022597   EA6050           CB-LCC-S
FPC 3
  PIC 0           REV 08   750-035293   EF3657           PF-1CGE-CFP
  PIC 1           REV 10   750-034624   BBAN4098         PF-12XGE-SFPP
FPC 5             REV 03   710-033871   BBAJ0768         T1600-FPC4-ES
  PIC 1           REV 03   750-034781   EE6655           PD-1CE-CFP-FPC4
FPC 6
  PIC 0           REV 10   750-034624   BBAN4109         PF-12XGE-SFPP
Fan Tray 0                                               FANTRAY-T-S
Fan Tray 1                                               FANTRAY-T4000-S
Fan Tray 2                                               FAN-REAR-TXP-LCC
```

show chassis hardware lcc (TX Matrix Router)

user@host> **show chassis hardware lcc 0**

```
lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                 65751            T640
Midplane          REV 03   710-005608   RA1408           T640 Backplane
FPM GBUS          REV 09   710-002901   RA2784           T640 FPM Board
FPM Display       REV 05   710-002897   RA2825           FPM Display
CIP               REV 06   710-002895   HT0684           T Series CIP
```



```
PEM 0             Rev 11   740-002595   PM18483          Power Entry Module
PEM 1             Rev 11   740-002595   qb13984          Power Entry Module
SCG 0             REV 11   710-003423   HT0022           T640 Sonet Clock Gen.
Routing Engine 0  REV 13   740-005022   210865700363     RE-3.0 (RE-600)
CB 0              REV 03   710-007655   HW1195           Control Board (CB-T)
FPC 1             REV 05   710-007527   HM3245           FPC Type 2
  CPU             REV 14   710-001726   HM1084           FPC CPU
  PIC 0           REV 02   750-007218   AZ1112           2x OC-12 ATM2 IQ, SMIR
  PIC 1           REV 02   750-007745   HG3462           4x OC-3 SONET, SMIR
  PIC 2           REV 14   750-001901   BA5390           4x OC-12 SONET, SMIR
  PIC 3           REV 09   750-008155   HS3012           2x G/E IQ, 1000 BASE
    SFP 0                  NON-JNPR     P1186TY          SFP-S
    SFP 1         REV 01   740-007326   P11WLTF          SFP-SX
  MMB 1           REV 02   710-005555   HL7514           MMB-288mbit
  PPB 0           REV 04   710-003758   HM4405           PPB Type 2
  PPB 1           REV 04   710-003758   AV1960           PPB Type 2
FPC 2             REV 08   710-010154   HZ3578           E-FPC Type 3
  CPU             REV 05   710-010169   HZ3219           FPC CPU-Enhanced
  PIC 0           REV 02   750-009567   HX2882           1x 10GE(LAN),XENPAK
    SFP 0         REV 01   740-009898   USC202U709       XENPAK-LR
  PIC 1           REV 03   750-003336   HJ9954           4x OC-48 SONET, SMSR
  PIC 2           REV 01   750-004535   HC0235           1x OC-192 SM SR1
  PIC 3           REV 07   750-007141   HX1699           10x 1GE(LAN), 1000 BASE

    SFP 0         REV 01   740-007326   2441042          SFP-SX
    SFP 1         REV 01   740-007326   2441027          SFP-SX
  MMB 0           REV 03   710-010171   HV2365           MMB-5M3-288mbit
  MMB 1           REV 03   710-010171   HZ3888           MMB-5M3-288mbit
SPMB 0            REV 09   710-003229   HW5245           T Series Switch CPU
SIB 3             REV 07   710-005781   HR5927           SIB-L8-F16
  B Board         REV 06   710-005782   HR5971           SIB-L8-F16 (B)
SIB 4             REV 07   710-005781   HR5903           SIB-L8-F16
  B Board         REV 06   710-005782   HZ5275           SIB-L8-F16 (B)
```

show chassis hardware scc (TX Matrix Router)

```
user@host> show chassis hardware scc
scc-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                                  TX Matrix
Midplane          REV 04   710-004396   RB0014           SCC Midplane
FPM GBUS          REV 04   710-004617   HW9141           SCC FPM Board
FPM Display       REV 04   710-004619   HS5950           SCC FPM
CIP 0             REV 01   710-010218   HV9151           SCC CIP
CIP 1             REV 01   710-010218   HV9152           SCC CIP
PEM 1             Rev 11   740-002595   QB13977          Power Entry Module
Routing Engine 0  REV 05   740-008883   P11123900153     RE-4.0 (RE-1600)
CB 0              REV 01   710-011709   HR5964           Control Board (CB-TX)
SPMB 0            REV 09   710-003229   HW5293           T Series Switch CPU
SIB 3
SIB 4             REV 01   710-005839   HW1177           SIB-S8-F16
  B Board         REV 01   710-005840   HW1202           SIB-S8-F16 (B)
```

show chassis hardware (T1600 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number    Description
```



```
Chassis                                  B2703             T1600
Midplane         REV 03   710-005608     RC4137            T640 Backplane
FPM GBUS         REV 10   710-002901     DT7062            T640 FPM Board
FPM Display      REV 05   710-002897     DS3067            FPM Display
CIP              REV 06   710-002895     DT3386            T-series CIP
PEM 0            Rev 07   740-017906     UA26344           Power Entry Module 3x80
PEM 1            Rev 18   740-002595     UF38441           Power Entry Module
SCG 0            REV 15   710-003423     DV0941            T640 Sonet Clock Gen.
Routing Engine 0 REV 08   740-014082     9009014502        RE-A-2000
Routing Engine 1 REV 07   740-014082     9009009591        RE-A-2000
CB 0             REV 05   710-007655     JA9360            Control Board (CB-T)
CB 1             REV 03   710-017707     DT3251            Control Board (CB-T)
FPC 0            REV 07   710-013558     DR4253            E2-FPC Type 2
  CPU            REV 05   710-013563     DS3902            FPC CPU-Enhanced
  PIC 0          REV 01   750-010618     CB5446            4x G/E SFP, 1000 BASE
    Xcvr 0       REV 01   740-011613     P9F11CW           SFP-SX
    Xcvr 1       REV 01   740-011613     P9F15C2           SFP-SX
    Xcvr 2       REV 01   740-011782     PB94K0L           SFP-SX
  PIC 1          REV 06   750-001900     HB6399            1x OC-48 SONET, SMSR
  PIC 2          REV 14   750-001901     AP1092            4x OC-12 SONET, SMIR
  PIC 3          REV 07   750-001900     AR8275            1x OC-48 SONET, SMSR
  MMB 1          REV 07   710-010171     DS1524            MMB-5M3-288mbit
FPC 1            REV 06   710-013553     DL9067            E2-FPC Type 1
  CPU            REV 04   710-013563     DM1685            FPC CPU-Enhanced
  PIC 0          REV 08   750-001072     AB1688            1x G/E, 1000 BASE-SX
  PIC 1          REV 10   750-012266     JX5519            4x 1GE(LAN), IQ2
    Xcvr 0       REV 01   740-011613     AM0812S8UK6       SFP-SX
    Xcvr 2       REV 01   740-011613     AM0812S8UK1       SFP-SX
    Xcvr 3       REV 01   740-011782     P8N1YHG           SFP-SX
  PIC 2          REV 22   750-005634     DP0083            1x CHOC12 IQ SONET, SMIR

  MMB 1          REV 07   710-008923     DN1862            MMB 3M 288-bit
FPC 2            REV 01   710-005548     HJ9899            FPC Type 3
  CPU            REV 06   710-001726     HC0586            FPC CPU
  PIC 0          REV 16   750-007141     NC9660            10x 1GE(LAN), 1000 BASE

    Xcvr 0       REV 01   740-011613     AM0812S8XAR       SFP-SX
    Xcvr 1       REV 01   740-011782     P920E7B           SFP-SX
    Xcvr 2       REV 01   740-011613     AM0812S8XAU       SFP-SX
    Xcvr 4       REV 01   740-011613     AM0812S8XAK       SFP-SX
    Xcvr 5       REV 01   740-011613     AM0812S8XAA       SFP-SX
    Xcvr 6       REV 01   740-011613     PAJ4NKY           SFP-SX
    Xcvr 7       REV 01   740-011613     AM0812S8UJW       SFP-SX
    Xcvr 8       REV 01   740-011782     PB81X89           SFP-SX
    Xcvr 9       REV 01   740-011613     AM0812S8UJX       SFP-SX
  PIC 1          REV 06   750-015217     DK3280            8x 1GE(TYPE3), IQ2
    Xcvr 0       REV 01   740-011782     P8P0A3T           SFP-SX
    Xcvr 1       REV 01   740-013111     5090002           SFP-T
    Xcvr 2       REV 01   740-011613     AM0814S93BQ       SFP-SX
    Xcvr 4                NON-JNPR       PDE0FAN           SFP-SX
    Xcvr 5       REV 01   740-011782     P8Q20XY           SFP-SX
    Xcvr 6       REV 01   740-011613     AM0812S8UJV       SFP-SX
    Xcvr 7       REV 01   740-011613     AM0812S8UP7       SFP-SX
  PIC 2          REV 05   750-004695     HT4383            1x Tunnel
  PIC 3          REV 17   750-009553     RL0204            4x OC-48 SONET
    Xcvr 0       REV 01   740-011785     PDS3T23           SFP-SR
    Xcvr 1       REV 01   740-011785     P6Q0F3E           SFP-SR
  MMB 0          REV 03   710-004047     HD5843            MMB-288mbit
  MMB 1          REV 03   710-004047     HE3208            MMB-288mbit
  PPB 0          REV 02   710-002845     HA4524            PPB Type 3
  PPB 1          REV 02   710-002845     HA4766            PPB Type 3
```



```
FPC 3            REV 01   710-010154   HR0863           E-FPC Type 3
  CPU            REV 01   710-010169   HN3422           FPC CPU-Enhanced
  PIC 0          REV 07   750-012793   WF5096           1x 10GE(LAN/WAN) IQ2
    Xcvr 0                NON-JNPR     M64294TP         XFP-10G-LR
  PIC 1          REV 25   750-007141   DV2127           10x 1GE(LAN), 1000 BASE

    Xcvr 0       REV 01   740-011613   PFA6LTJ          SFP-SX
    Xcvr 1       REV 01   740-011782   P9P0XV4          SFP-SX
    Xcvr 2       REV 01   740-011782   P9M0TNX          SFP-SX
    Xcvr 4       REV 01   740-011782   P9B0TTP          SFP-SX
    Xcvr 5                NON-JNPR     PBS4LED          SFP-SX
  PIC 2          REV 17   750-009553   RL0212           4x OC-48 SONET
    Xcvr 0       REV 01   740-011785   PDS3T8G          SFP-SR
  PIC 3          REV 32   750-003700   DL1279           1x OC-192 12xMM VSR
  MMB 0          REV 01   710-010171   HR0821           MMB-288mbit
  MMB 1          REV 01   710-010171   HR0818           MMB-288mbit
FPC 4            REV 16   710-013037   EB4919           FPC Type 4-ES
  CPU            REV 09   710-016744   BBAA4382         ST-PMB2
  PIC 0          REV 03   711-029996   EB1569           100GE
  PIC 1          REV 05   711-029999   EB9983           100GE CFP
    Xcvr 0       REV 0    740-032210   J10G80746        CFP-100G-LR4
  BRIDGE 0       REV 02   711-029995   EB2235           100GE Bridge Board
  MMB 0          REV 04   710-025563   BBAA7112         ST-MMB2
  MMB 1          REV 04   710-025563   BBAA7149         ST-MMB2
FPC 5            REV 02   710-013037   DE3407           FPC Type 4-ES
  CPU            REV 04   710-016744   DA2124           ST-PMB2
  PIC 0          REV 16   750-012518   DF2554           4x OC-192 SONET XFP
    Xcvr 0       REV 01   740-014279   AA0745N1FX8      XFP-OC192-SR
    Xcvr 1       REV 01   740-014279   AA0748N1HN5      XFP-OC192-SR
    Xcvr 2       REV 01   740-014279   AA0748N1HT6      XFP-OC192-SR
    Xcvr 3       REV 01   740-014279   AA0744N1EC9      XFP-OC192-SR
  PIC 1          REV 01   750-010850   JA0329           1x OC-768 SONET SR
  MMB 0          REV 04   710-016036   DE9577           ST-MMB2
  MMB 1          REV 04   710-016036   DK4060           ST-MMB2
FPC 6            REV 14   710-013037   DV1431           FPC Type 4-ES
  CPU            REV 09   710-016744   DT9020           ST-PMB2
  PIC 0          REV 11   750-017405   DM6261           4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 01   740-014289   C701XU05Q        XFP-10G-SR
    Xcvr 1       REV 01   740-014279   AA0748N1HPT      XFP-10G-LR
    Xcvr 2       REV 01   740-014289   T08E19189        XFP-10G-SR
    Xcvr 3       REV 01   740-014289   C715XU058        XFP-10G-SR
  PIC 1          REV 13   750-017405   DP8772           4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 02   740-011571   C850XJ037        XFP-10G-SR
    Xcvr 1       REV 02   740-014289   C839XU0L9        XFP-10G-SR
    Xcvr 2       REV 02   740-014289   C834XU05A        XFP-10G-SR
    Xcvr 3       REV 02   740-014289   C810XU0CE        XFP-10G-SR
  MMB 0          REV 01   710-025563   DT8454           ST-MMB2
  MMB 1          REV 01   710-025563   DT8366           ST-MMB2
FPC 7            REV 09   710-007529   HZ7624           FPC Type 3
  CPU            REV 15   710-001726   HZ1413           FPC CPU
  PIC 0          REV 10   750-012793   DM5627           1x 10GE(LAN/WAN) IQ2
    Xcvr 0       REV 02   740-011571   C831XJ062        XFP-10G-SR
  PIC 1          REV 01   750-015217   JT6762           8x 1GE(TYPE3), IQ2
    Xcvr 0       REV 01   740-011782   P8Q25JU          SFP-SX
    Xcvr 1       REV 01   740-011782   P9B0U0K          SFP-SX
  PIC 2          REV 01   750-015217   JS4268           8x 1GE(TYPE3), IQ2
    Xcvr 0       REV 01   740-011613   AM0812S8XBZ      SFP-SX
    Xcvr 1       REV 01   740-011613   AM0812S8XAP      SFP-SX
    Xcvr 2       REV 01   740-011613   AM0812S8XBY      SFP-SX
    Xcvr 3       REV 01   740-011613   AM0812S8XBX      SFP-SX
    Xcvr 4       REV 01   740-011613   P9F1652          SFP-SX
```

```
      Xcvr 5        REV 01   740-011782   P8Q21YC          SFP-SX
      Xcvr 6        REV 01   740-011782   P8Q27HQ          SFP-SX
      Xcvr 7        REV 01   740-011613   P8E2SSU          SFP-SX
    PIC 3           REV 15   750-009450   NB6790           1x OC-192 SM SR2
    MMB 0           REV 03   710-005555   HZ3450           MMB-288mbit
    MMB 1           REV 03   710-005555   HZ3415           MMB-288mbit
    PPB 0           REV 04   710-002845   HP0887           PPB Type 3
    PPB 1           REV 04   710-002845   HW5255           PPB Type 3
  SPMB 0            REV 10   710-003229   HX3699           T-series Switch CPU
  SPMB 1            REV 12   710-003229   DT3091           T-series Switch CPU
  SIB 0             REV 07   710-013074   DS4747           SIB-I8-SF
  SIB 1             REV 07   710-013074   DS4942           SIB-I8-SF
  SIB 2             REV 07   710-013074   DS4965           SIB-I8-SF
  SIB 3             REV 07   710-013074   DS4990           SIB-I8-SF
  SIB 4             REV 07   710-013074   DS4944           SIB-I8-SF
  Fan Tray 0                                               Front Top Fan Tray
  Fan Tray 1                                               Front Bottom Fan Tray
  Fan Tray 2                                               Rear Fan Tray -- Rev 2
```

show chassis hardware (TX Matrix Plus Router)

```
user@host> show chassis hardware
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number    Description
Chassis                                 JN113186EAHB     TXP
Midplane          REV 05   710-022574   TS3822           SFC Midplane
FPM Display       REV 03   710-024027   DW4701           TXP FPM Display
CIP 0             REV 05   710-023792   DW7998           TXP CIP
CIP 1             REV 05   710-023792   DW7999           TXP CIP
PEM 0             Rev 04   740-027463   UM26367          Power Entry Module
PEM 1             Rev 04   740-027463   UM26346          Power Entry Module
Routing Engine 0  REV 06   740-026942   737A-1081        RE-DUO-2600
Routing Engine 1  REV 06   740-026942   737A-1043        RE-DUO-2600
CB 0              REV 05   710-022606   DW4435           SFC Control Board
CB 1              REV 09   710-022606   DW6100           SFC Control Board
SPMB 0                     BUILTIN                       SFC Switch CPU
SPMB 1                     BUILTIN                       SFC Switch CPU
SIB F13 0         REV 04   750-024564   DW5764           F13 SIB
  B Board         REV 03   710-023431   DW9053           F13 SIB Mezz
SIB F13 3         REV 04   750-024564   DW5785           F13 SIB
  B Board         REV 03   710-023431   DW9030           F13 SIB Mezz
SIB F13 6
SIB F13 8         REV 04   750-024564   DW5752           F13 SIB
  B Board         REV 03   710-023431   DW9051           F13 SIB Mezz
SIB F13 11        REV 04   750-024564   DW5782           F13 SIB
  B Board         REV 03   710-023431   DW9058           F13 SIB Mezz
SIB F13 12        REV 03   750-024564   DT9466           F13 SIB
  B Board         REV 02   710-023431   DT6556           F13 SIB Mezz
SIB F2S 0/0       REV 05   710-022603   DW7898           F2S SIB
  B Board         REV 05   710-023787   DW7625           F2S SIB Mezz
SIB F2S 0/2       REV 05   710-022603   DW7811           F2S SIB
  B Board         REV 05   710-023787   DW7550           F2S SIB Mezz
SIB F2S 0/4       REV 04   710-022603   DW4873           F2S SIB
  B Board         REV 05   710-023787   DW8509           F2S SIB Mezz
SIB F2S 0/6       REV 04   710-022603   DW4867           F2S SIB
  B Board         REV 05   710-023787   DW8472           F2S SIB Mezz
SIB F2S 1/0       REV 04   710-022603   DW4871           F2S SIB
  B Board         REV 05   710-023787   DW8497           F2S SIB Mezz
SIB F2S 1/2       REV 05   710-022603   DW7868           F2S SIB
```



```
  B Board        REV 05   710-023787   DW7551           F2S SIB Mezz
SIB F2S 1/4      REV 04   710-022603   DW4854           F2S SIB
  B Board        REV 05   710-023787   DW8496           F2S SIB Mezz
SIB F2S 1/6      REV 05   710-022603   DW7889           F2S SIB
  B Board        REV 05   710-023787   DW7496           F2S SIB Mezz
SIB F2S 2/0      REV 04   710-022603   DW4852           F2S SIB
  B Board        REV 05   710-023787   DW8498           F2S SIB Mezz
SIB F2S 2/2      REV 04   710-022603   DW4845           F2S SIB
  B Board        REV 05   710-023787   DW8457           F2S SIB Mezz
SIB F2S 2/4      REV 05   710-022603   DW7802           F2S SIB
  B Board        REV 05   710-023787   DW7562           F2S SIB Mezz
SIB F2S 2/6      REV 04   710-022603   DW4822           F2S SIB
  B Board        REV 05   710-023787   DW8467           F2S SIB Mezz
SIB F2S 3/0      REV 05   710-022603   DW7815           F2S SIB
  B Board        REV 05   710-023787   DW7518           F2S SIB Mezz
SIB F2S 3/2      REV 03   710-022603   DV0068           F2S SIB
  B Board        REV 03   710-023787   DT9974           F2S SIB Mezz
SIB F2S 3/4      REV 05   710-022603   DW7874           F2S SIB
  B Board        REV 05   710-023787   DW7601           F2S SIB Mezz
SIB F2S 3/6      REV 03   710-022603   DV0033           F2S SIB
  B Board        REV 03   710-023787   DT9969           F2S SIB Mezz
SIB F2S 4/0      REV 03   710-022603   DV0043           F2S SIB
  B Board        REV 03   710-023787   DT9948           F2S SIB Mezz
SIB F2S 4/2      REV 05   710-022603   DW5446           F2S SIB
  B Board        REV 05   710-023787   DW7611           F2S SIB Mezz
SIB F2S 4/4      REV 04   710-022603   DW4826           F2S SIB
  B Board        REV 05   710-023787   DW8458           F2S SIB Mezz
SIB F2S 4/6      REV 03   710-022603   DV0026           F2S SIB
  B Board        REV 03   710-023787   DT9963           F2S SIB Mezz
Fan Tray 0       REV 02   760-024497   DR8290           Front Fan Tray
Fan Tray 1       REV 02   760-024497   DR8293           Front Fan Tray
Fan Tray 2       REV 05   760-024502   DR8280           Rear Fan Tray
Fan Tray 3
Fan Tray 4       REV 05   760-024502   DR8276           Rear Fan Tray
Fan Tray 5       REV 02   760-024502   DP5643           Rear Fan Tray

lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN11036F8AHA     T1600
Midplane         REV 03   710-017247   RC3799           T-series Backplane
FPM GBUS         REV 10   710-002901   DP7009           T640 FPM Board
FPM Display      REV 01   710-021387   DN7026           T1600 FPM Display
CIP              REV 06   710-002895   DP6024           T-series CIP
PEM 1           Rev 02   740-023211   WA50019          Power Entry Module 4x60A
SCG 0            REV 15   710-003423   DR6757           T640 Sonet Clock Gen.
SCG 1            REV 15   710-003423   DS2225           T640 Sonet Clock Gen.
Routing Engine 0 REV 01   740-026941   737F-1040        RE-DUO-1800
Routing Engine 1 REV 01   740-026941   737F-1016        RE-DUO-1800
CB 0             REV 06   710-022597   DX4011           LCC Control Board
CB 1             REV 06   710-022597   DX4017           LCC Control Board
FPC 1            REV 07   710-013035   DN5847           FPC Type 3-ES
  CPU            REV 08   710-016744   DP2570           ST-PMB2
  PIC 0          REV 05   750-015217   DB0418           8x 1GE(TYPE3), IQ2
    Xcvr 0       REV 01   740-011782   P8Q27ZG          SFP-SX
    Xcvr 1                NON-JNPR     PDA1U0D          SFP-SX
    Xcvr 2       REV 01   740-011613   P9F1ALW          SFP-SX
    Xcvr 3       REV 01   740-011782   PBA403V          SFP-SX
    Xcvr 4                NON-JNPR     PDE09DP          SFP-SX
    Xcvr 5       REV 01   740-011782   PCH2P4K          SFP-SX
```

```
      Xcvr 6        REV 01   740-011782   PB94K0F          SFP-SX
      Xcvr 7        REV 01   740-011782   PBA2R2A          SFP-SX
    PIC 1           REV 03   750-004424   HJ4020           1x 10GE(LAN),DWDM
    PIC 2           REV 01   750-003336   HG6073           4x OC-48 SONET, SMSR
    MMB 0           REV 04   710-016036   DP3401           ST-MMB2
  FPC 3             REV 12   710-013037   DR1169           FPC Type 4-ES
    CPU             REV 08   710-016744   DP9429           ST-PMB2
    PIC 0           REV 02   750-010850   JA0332           1x OC-768 SONET SR
    MMB 0           REV 04   710-016036   DR0628           ST-MMB2
    MMB 1           REV 04   710-016036   DR0592           ST-MMB2
  FPC 4             REV 05   710-021534   DR7350           FPC Type 1-ES
    CPU             REV 08   710-016744   DP8096           ST-PMB2
    PIC 0           REV 04   750-014627   DP9171           4x OC-3 1x OC-12 SFP
      Xcvr 0        REV 02   740-011615   PDE2RVR          SFP-SR
    PIC 1          REV 22   750-005634   DS5815          1x CHOC12 IQ SONET, SMIR

    PIC 2           REV 09   750-002911   CF4539           4x F/E, 100 BASE-TX
    PIC 3           REV 08   750-021652   DR2827           1x CHOC12 IQE SONET
      Xcvr 0                 NON-JNPR     8                UNKNOWN
    MMB 0           REV 04   710-016036   DR0809           ST-MMB2
  FPC 5             REV 07   710-007529   HS5608           FPC Type 3
    CPU             REV 15   710-001726   HX4351           FPC CPU
    PIC 0           REV 14   750-009567   WJ8961           1x 10GE(LAN),XENPAK
      Xcvr 0        REV 01   740-013170   J05K05961        XENPAK-LR
    PIC 1           REV 16   750-007141   JJ8146           10x 1GE(LAN), 1000 BASE

      Xcvr 1        REV 01   740-011613   P9F117T          SFP-SX
      Xcvr 2        REV 01   740-011782   PBA2VCL          SFP-SX
      Xcvr 3        REV 01   740-011782   PB83DRB          SFP-SX
      Xcvr 4        REV 01   740-011613   AM0812S8UP8      SFP-SX
    PIC 2           REV 12   750-009567   WF3566           1x 10GE(LAN),XENPAK
      Xcvr 0        REV 02   740-013170   T07C94489        XENPAK-LR
    MMB 0           REV 03   710-005555   HZ1907           MMB-288mbit
    MMB 1           REV 03   710-005555   HW5283           MMB-288mbit
    PPB 0           REV 04   710-002845   HZ7717           PPB Type 3
    PPB 1           REV 04   710-002845   HS0110           PPB Type 3
  FPC 6             REV 07   710-013035   DP7486           FPC Type 3-ES
    CPU             REV 08   710-016744   DP2545           ST-PMB2
    PIC 0           REV 09   750-009567   NE6323           1x 10GE(LAN),XENPAK
      Xcvr 0        REV 02   740-013170   T09C71959        XENPAK-LR
    PIC 1           REV 06   750-015217   DN4775           8x 1GE(TYPE3), IQ2
      Xcvr 0        REV 01   740-011782   P7E0T6M          SFP-SX
      Xcvr 1        REV 01   740-011613   AM0812S8XAY      SFP-SX
      Xcvr 2        REV 01   740-011782   P7E0T6J          SFP-SX
      Xcvr 3        REV 01   740-011782   PCH2P7D          SFP-SX
      Xcvr 4        REV 01   740-011782   P9B0QYT          SFP-SX
      Xcvr 5        REV 01   740-011613   AM0812S8WQJ      SFP-SX
      Xcvr 6        REV 02   740-013111   9301220          SFP-T
      Xcvr 7        REV 01   740-011782   P9B0TZ5          SFP-SX
    PIC 2           REV 06   750-015217   DM6747           8x 1GE(TYPE3), IQ2
      Xcvr 0        REV 01   740-011613   PAP0ZB2          SFP-SX
      Xcvr 1        REV 01   740-013111   70191002         SFP-T
      Xcvr 6        REV 01   740-011782   PBA29H8          SFP-SX
      Xcvr 7        REV 01   740-011613   AM0812S8WQG      SFP-SX
    MMB 0           REV 04   710-016036   DP3238           ST-MMB2
  FPC 7             REV 03   710-021540   DV3154           FPC Type 2-ES
    CPU             REV 09   710-016744   DT9053           ST-PMB2
    PIC 0           REV 13   750-001901   HB4225           4x OC-12 SONET, SMIR
    PIC 1           REV 05   750-001900   AD3644           1x OC-48 SONET, SMSR
    PIC 2           REV 10   750-008155   HV0335           2x G/E IQ, 1000 BASE
      Xcvr 0        REV 01   740-011782   PCH2UKF          SFP-SX
```



```
    Xcvr 1        REV 01   740-011782   PCH2V19           SFP-SX
  PIC 3           REV 03   750-014638   JS9493            1x OC-48-12-3 SFP
    Xcvr 0        REV 01   740-011785   P6Q0ENK           SFP-SR
  MMB 0           REV 05   710-016036   DP3323            ST-MMB2
SPMB 0            REV 04   710-023321   DX3004            LCC Switch CPU
SPMB 1            REV 04   710-023321   DX3009            LCC Switch CPU
SIB 0             REV 07   710-022594   DW4195            LCC SIB
  B Board         REV 07   710-023185   DW3930            LCC SIB Mezz
SIB 1             REV 07   710-022594   DW4179            LCC SIB
  B Board         REV 07   710-023185   DW3919            LCC SIB Mezz
SIB 2
SIB 3             REV 06   710-022594   DT8251            LCC SIB
  B Board         REV 06   710-023185   DT5792            LCC SIB Mezz
SIB 4             REV 08   710-022594   DW8014            LCC SIB
  B Board         REV 07   710-023185   DW3917            LCC SIB Mezz
Fan Tray 0                                                Front Top Fan Tray
Fan Tray 1                                                Front Bottom Fan Tray
Fan Tray 2                                                Rear Fan Tray -- Rev 3

lcc1-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number     Description
Chassis                                 JN1102270AHA      T1600
Midplane          REV 04   710-017247   RC5358            T-series Backplane
FPM GBUS          REV 10   710-002901   DS3443            T640 FPM Board
FPM Display       REV 01   710-021387   DS6411            T1600 FPM Display
CIP               REV 06   710-002895   DS4235            T-series CIP
PEM 0            Rev 02   740-023211   VM82438           Power Entry Module 4x60A
SCG 0             REV 15   710-003423   DS6649            T640 Sonet Clock Gen.
SCG 1             REV 15   710-003423   DR6775            T640 Sonet Clock Gen.
Routing Engine 0  REV 01   740-026941   737F-1083         RE-DUO-1800
Routing Engine 1  REV 01   740-026941   737F-1104         RE-DUO-1800
CB 0              REV 06   710-022597   DW8542            LCC Control Board
CB 1              REV 06   710-022597   DW8530            LCC Control Board
FPC 0             REV 02   710-010845   JE2392            FPC Type 4
  CPU             REV 02   710-011481   JF6820            FPC CPU-Enhanced
  PIC 0           REV 11   750-017405   DP7259            4x 10GE (LAN/WAN) XFP
    Xcvr 0        REV 01   740-014279   AA0741N1C8T       XFP-10G-LR
    Xcvr 1        REV 01   740-014279   AA0746N1GAM       XFP-10G-LR
    Xcvr 2        REV 01   740-014279   AA0747N1H0B       XFP-10G-LR
    Xcvr 3        REV 01   740-014279   AA0748N1HZ5       XFP-10G-LR
  MMB 0           REV 03   710-010842   HY7601            ST-MMB
FPC 1             REV 16   710-013037   BBAA7398          FPC Type 4-ES
  CPU             REV 09   710-016744   BBAA2329          ST-PMB2
  PIC 0           REV 03   711-029996   EB1575            100GE
  PIC 1           REV 06   750-034781   EB9980            100GE CFP
  MMB 0           REV 04   710-025563   BBAA5325          ST-MMB2
  MMB 1           REV 04   710-025563   BBAA5444          ST-MMB2
FPC 2             REV 16   710-013037   BBAA7185          FPC Type 4-ES
  CPU             REV 09   710-016744   BBAA3522          ST-PMB2
  PIC 0           REV 03   711-029996   EB1557            100GE
  PIC 1           REV 05   750-034781   EB4660            100GE CFP
    Xcvr 0        REV 0    740-032210   J10F73666         CFP-100G-LR4
  BRIDGE 0        REV 02   711-029995   EB2237            100GE Bridge Board
  MMB 0           REV 04   710-025563   BBAA5347          ST-MMB2
  MMB 1           REV 04   710-025563   BBAA5401          ST-MMB2
FPC 3             REV 10   710-021534   DZ0941            FPC Type 1-ES
  CPU             REV 09   710-016744   DY6364            ST-PMB2
  PIC 0           REV 13   750-012266   DK9192            4x 1GE(LAN), IQ2
    Xcvr 0        REV 01   740-011613   AM0812S8WVD       SFP-SX
```

```
    Xcvr 1                NON-JNPR    PDD63Q4          SFP-SX
    Xcvr 2                NON-JNPR    PDE4G54          SFP-SX
    Xcvr 3                NON-JNPR    PD40MAG          SFP-SX
  PIC 1          REV 01   750-007641  HJ2003           1x G/E IQ, 1000 BASE
    Xcvr 0       REV 01   740-011613  AM0812S8WVG      SFP-SX
  PIC 3          REV 17   750-007444  JB6873           1x CHSTM1 IQ SDH, SMIR
  MMB 0          REV 04   710-025563  DZ0281           ST-MMB2
FPC 4            REV 06   710-013035  DK0614           FPC Type 3-ES
  CPU            REV 07   710-016744  DK1616           ST-PMB2
  PIC 0          REV 22   750-007141  DM1870           10x 1GE(LAN), 1000 BASE

    Xcvr 0       REV 01   740-011782  PCL3UKW          SFP-SX
    Xcvr 1       REV 01   740-011782  P7E0T73          SFP-SX
    Xcvr 2       REV 01   740-007326  P4T0WLR          SFP-SX
    Xcvr 3       REV 01   740-011782  PAR1LRL          SFP-SX
    Xcvr 4       REV 01   740-011782  P9M0U3Z          SFP-SX
    Xcvr 5       REV 01   740-011782  P9M0U0C          SFP-SX
    Xcvr 6       REV 01   740-011782  P9M0TLG          SFP-SX
    Xcvr 7       REV 01   740-011782  P9M0U0F          SFP-SX
    Xcvr 8       REV 01   740-011613  PFA6LAP          SFP-SX
    Xcvr 9       REV 01   740-011782  PCH2P0U          SFP-SX
  PIC 1          REV 16   750-009450  CV2565           1x OC-192 SM SR2
  PIC 2          REV 05   750-004424  HH3057           1x 10GE(LAN),10GBASE-LR

  PIC 3          REV 12   750-013423  DP0403           MultiServices 500
  MMB 0          REV 04   710-016036  DK1988           ST-MMB2
FPC 5            REV 07   710-013560  DR0004           E2-FPC Type 3
  CPU            REV 05   710-013563  DR0089           FPC CPU-Enhanced
  PIC 0          REV 11   750-012793  DR6107           1x 10GE(LAN/WAN) IQ2
    Xcvr 0       REV 01   740-014289  C743XU074        XFP-10G-SR
  PIC 1          REV 01   750-004695  HD5980           1x Tunnel
  PIC 2          REV 32   750-003700  DL3770           1x OC-192 12xMM VSR
  PIC 3          REV 12   750-009553  WB8901           4x OC-48 SONET
    Xcvr 0       REV 01   740-011785  P9D1GTQ          SFP-SR
    Xcvr 1       REV 01   740-011785  PDS0MMB          SFP-SR
    Xcvr 3       REV 01   740-011785  PDE1KXP          SFP-SR
  MMB 0          REV 07   710-010171  DP7374           MMB-5M3-288mbit
  MMB 1          REV 07   710-010171  DP7404           MMB-5M3-288mbit
FPC 6            REV 07   710-013035  DM0994           FPC Type 3-ES
  CPU            REV 07   710-016744  DM3651           ST-PMB2
  PIC 0          REV 07   750-015217  DN4743           8x 1GE(TYPE3), IQ2
    Xcvr 3       REV 01   740-011613  AM0812S8XB0      SFP-SX
    Xcvr 4       REV 01   740-011782  PB829RB          SFP-SX
    Xcvr 5       REV 01   740-011782  P8J1SYX          SFP-SX
  PIC 1          REV 03   750-003336  HJ9954           4x OC-48 SONET, SMSR
  PIC 3          REV 02   750-012793  JM7665           1x 10GE(LAN/WAN) IQ2
  MMB 0          REV 04   710-016036  DN6913           ST-MMB2
FPC 7            REV 08   710-010845  JM3958           FPC Type 4
  CPU            REV 04   710-011481  JK3669           FPC CPU-Enhanced
  PIC 0          REV 11   750-017405  DP8837           4x 10GE (LAN/WAN) XFP
    Xcvr 1       REV 01   740-014279  753019A00277     XFP-10G-LR
    Xcvr 2       REV 02   740-011571  C850XJ00P        XFP-10G-SR
    Xcvr 3       REV 01   740-014279  AA0813N1RTG      XFP-10G-LR
  MMB 0          REV 04   710-010842  JN1971           ST-MMB
SPMB 0           REV 04   710-023321  DW3629           LCC Switch CPU
SPMB 1           REV 04   710-023321  DW3621           LCC Switch CPU
SIB 0            REV 07   710-022594  DW4200           LCC SIB
  B Board        REV 07   710-023185  DW3932           LCC SIB Mezz
SIB 1            REV 07   710-022594  DW4193           LCC SIB
  B Board        REV 07   710-023185  DW3904           LCC SIB Mezz
SIB 2
```

```
SIB 3            REV 07   710-022594   DW4210           LCC SIB
  B Board        REV 06   710-023185   DT5780           LCC SIB Mezz
SIB 4            REV 08   710-022594   DW8019           LCC SIB
  B Board        REV 06   710-023185   DT5795           LCC SIB Mezz
Fan Tray 0                                              Front Top Fan Tray
Fan Tray 1                                              Front Bottom Fan Tray
Fan Tray 2                                              Rear Fan Tray -- Rev 3
```

show chassis hardware sfc (TX Matrix Plus Router)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN112F007AHB     TXP
Midplane         REV 05   710-022574   TS4027           SFC Midplane
FPM Display      REV 03   710-024027   DX0282           TXP FPM Display
CIP 0            REV 04   710-023792   DW4889           TXP CIP
CIP 1            REV 04   710-023792   DW4887           TXP CIP
PEM 0            Rev 07   740-027463   UM26368          Power Entry Module
Routing Engine 0 REV 01   740-026942   737A-1064        SFC RE
Routing Engine 1 REV 01   740-026942   737A-1082        SFC RE
CB 0             REV 09   710-022606   DW6099           SFC Control Board
CB 1             REV 09   710-022606   DW6096           SFC Control Board
SPMB 0                    BUILTIN                       SFC Switch CPU
SPMB 1                    BUILTIN                       SFC Switch CPU
SIB F13 0        REV 04   710-022600   DX0841           F13 SIB
  B Board        REV 03   710-023431   DX0966           F13 SIB Mezz
SIB F13 1        REV 04   750-024564   DW5776           F13 SIB
  B Board        REV 03   710-023431   DW9028           F13 SIB
SIB F13 3        REV 04   750-024564   DW5762           F13 SIB
  B Board        REV 03   710-023431   DW9059           F13 SIB
SIB F13 4        REV 04   750-024564   DW5797           F13 SIB
  B Board        REV 03   710-023431   DW9041           F13 SIB
SIB F13 6        REV 04   750-024564   DW5770           F13 SIB
  B Board        REV 03   710-023431   DW9079           F13 SIB Mezz
SIB F13 7        REV 04   750-024564   DW5758           F13 SIB
  B Board        REV 03   710-023431   DW9047           F13 SIB
SIB F13 8        REV 04   750-024564   DW5761           F13 SIB
  B Board        REV 03   710-023431   DW9043           F13 SIB Mezz
SIB F13 9        REV 04   750-024564   DW5754           F13 SIB
  B Board        REV 03   710-023431   DW9078           F13 SIB Mezz
SIB F13 11       REV 04   710-022600   DX0826           F13 SIB
  B Board        REV 03   710-023431   DX0967           F13 SIB Mezz
SIB F13 12       REV 04   750-024564   DW5794           F13 SIB
  B Board        REV 03   710-023431   DW9044           F13 SIB Mezz
SIB F2S 0/0      REV 05   710-022603   DW7897           F2S SIB
  B Board        REV 05   710-023787   DW7657           NEO PMB
SIB F2S 0/2      REV 05   710-022603   DW7833           F2S SIB
  B Board        REV 05   710-023787   DW7526           NEO PMB
SIB F2S 0/4      REV 05   710-022603   DW7875           F2S SIB
  B Board        REV 05   710-023787   DW7588           NEO PMB
SIB F2S 0/6      REV 05   710-022603   DW7860           F2S SIB
  B Board        REV 05   710-023787   DW7589           NEO PMB
SIB F2S 1/0      REV 04   710-022603   DW4820           F2S SIB
  B Board        REV 05   710-023787   DW8510           NEO PMB
SIB F2S 1/2      REV 05   710-022603   DW7849           F2S SIB
  B Board        REV 05   710-023787   DW7525           NEO PMB
SIB F2S 1/4      REV 05   710-022603   DW7927           F2S SIB
  B Board        REV 05   710-023787   DW7556           F2S SIB Mezz
```

```
SIB F2S 1/6       REV 05   710-022603   DW7866            F2S SIB
  B Board         REV 05   710-023787   DW7651            NEO PMB
SIB F2S 2/0       REV 05   710-022603   DW7880            F2S SIB
  B Board         REV 05   710-023787   DW7523            NEO PMB
SIB F2S 2/2       REV 05   710-022603   DW7895            F2S SIB
  B Board         REV 05   710-023787   DW7591            NEO PMB
SIB F2S 2/4       REV 05   710-022603   DW7907            F2S SIB
  B Board         REV 05   710-023787   DW7590            NEO PMB
SIB F2S 2/6       REV 05   710-022603   DW7785            F2S SIB
  B Board         REV 05   710-023787   DW7524            NEO PMB
SIB F2S 3/0       REV 05   710-022603   DW7782            F2S SIB
  B Board         REV 05   710-023787   DW7634            NEO PMB
SIB F2S 3/2       REV 05   710-022603   DW7793            F2S SIB
  B Board         REV 05   710-023787   DW7548            NEO PMB
SIB F2S 3/4       REV 05   710-022603   DW7779            F2S SIB
  B Board         REV 05   710-023787   DW7587            NEO PMB
SIB F2S 3/6       REV 05   710-022603   DW7930            F2S SIB
  B Board         REV 05   710-023787   DW7505            NEO PMB
SIB F2S 4/0       REV 05   710-022603   DW7867            F2S SIB
  B Board         REV 05   710-023787   DW7656            NEO PMB
SIB F2S 4/2       REV 05   710-022603   DW7917            F2S SIB
  B Board         REV 05   710-023787   DW7640            NEO PMB
SIB F2S 4/4       REV 05   710-022603   DW7929            F2S SIB
  B Board         REV 05   710-023787   DW7643            NEO PMB
SIB F2S 4/6       REV 05   710-022603   DW7870            F2S SIB
  B Board         REV 05   710-023787   DW7635            NEO PMB
Fan Tray 0        REV 06   760-024497   DV7831            Front Fan Tray
Fan Tray 1        REV 06   760-024497   DV9614            Front Fan Tray
Fan Tray 2        REV 06   760-024502   DV9618            Rear Fan Tray
Fan Tray 3        REV 06   760-024502   DV9616            Rear Fan Tray
Fan Tray 4        REV 06   760-024502   DV7807            Rear Fan Tray
Fan Tray 5        REV 06   760-024502   DV7828            Rear Fan Tray
```

show chassis hardware extensive (TX Matrix Plus Router)

```
user@host> show chassis hardware extensive
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number     Description
Chassis                                 JN112F007AHB      TXP
 Jedec Code:    0x7fb0            EEPROM Version:    0x02
                                  S/N:               JN112F007AHB
 Assembly ID:   0x052c            Assembly Version:  00.00
 Date:          00-00-0000        Assembly Flags:    0x00
 ID: TXP
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 2c 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 4a 4e 31 31 32 46 30 30 37 41 48 42 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 05   710-022574   TS4027            SFC Midplane
 Jedec Code:    0x7fb0            EEPROM Version:    0x01
 P/N:           710-022574        S/N:               TS4027
 Assembly ID:   0x0962            Assembly Version:  01.05
```



```
 Date:          03-23-2009        Assembly Flags:    0x00
 Version:       REV 05
 ID: SFC Midplane
 Board Information Record:
  Address 0x00: ad 01 ff ff 00 1d b5 14 00 00 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 62 01 05 52 45 56 20 30 35 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 32 32 35 37 34 00 00
  Address 0x20: 53 2f 4e 20 54 53 34 30 32 37 00 00 00 17 03 07
  Address 0x30: d9 ff ff ff ad 01 ff ff 00 1d b5 14 00 00 ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Display       REV 03   710-024027   DX0282          TXP FPM Display
 Jedec Code:    0x7fb0            EEPROM Version:    0x01
 P/N:           710-024027        S/N:               DX0282
 Assembly ID:   0x096c            Assembly Version:  01.03
 Date:          02-10-2009        Assembly Flags:    0x00
 Version:       REV 03
 ID: TXP FPM Display              FRU Model Number:  CRAFT-TXP
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 09 6c 01 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 30 32 37 00 00
  Address 0x20: 53 2f 4e 20 44 58 30 32 38 32 00 00 00 0a 02 07
  Address 0x30: d9 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
  Address 0x50: 52 41 46 54 2d 54 58 50 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CIP 0             REV 04   710-023792   DW4889          TXP CIP
 Jedec Code:    0x7fb0            EEPROM Version:    0x01
 P/N:           710-023792        S/N:               DW4889
 Assembly ID:   0x0969            Assembly Version:  01.04
 Date:          01-26-2009        Assembly Flags:    0x00
 Version:       REV 04
 ID: TXP CIP                      FRU Model Number:  CIP-TXP
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware clei-models (TX Matrix Plus Router)

```
user@host> show chassis hardware clei-models
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  CLEI code         FRU model number
Midplane          REV 05   710-022574                     CHAS-BP-TXP-S
FPM Display       REV 03   710-024027                     CRAFT-TXP-S
CIP 0             REV 05   710-023792                     CIP-TXP-S
CIP 1             REV 05   710-023792                     CIP-TXP-S
PEM 0             Rev 04   740-027463   IPUPAFGKTA        PWR-TXP-7-60-DC
PEM 1             Rev 04   740-027463   IPUPAFGKTA        PWR-TXP-7-60-DC
Routing Engine 0  REV 06   740-026942                     RE-DUO-C2600-16G-S
Routing Engine 1  REV 06   740-026942                     RE-DUO-C2600-16G-S
CB 0              REV 05   710-022606                     CB-TXP-S
CB 1              REV 09   710-022606                     CB-TXP-S
SIB F13 0         REV 04   750-024564                     SIB-TXP-F13
SIB F13 3         REV 04   750-024564                     SIB-TXP-F13
```



```
SIB F13 8        REV 04   750-024564                      SIB-TXP-F13
SIB F13 11       REV 04   750-024564                      SIB-TXP-F13
SIB F13 12       REV 03   750-024564                      SIB-TXP-F13
SIB F2S 0/0      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 0/2      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 0/4      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 0/6      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 1/0      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 1/2      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 1/4      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 1/6      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 2/0      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 2/2      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 2/4      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 2/6      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 3/0      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 3/2      REV 03   710-022603                      SIB-TXP-F2S-S
SIB F2S 3/4      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 3/6      REV 03   710-022603                      SIB-TXP-F2S-S
SIB F2S 4/0      REV 03   710-022603                      SIB-TXP-F2S-S
SIB F2S 4/2      REV 05   710-022603                      SIB-TXP-F2S-S
SIB F2S 4/4      REV 04   710-022603                      SIB-TXP-F2S-S
SIB F2S 4/6      REV 03   710-022603                      SIB-TXP-F2S-S
Fan Tray 0       REV 02   760-024497                      FANTRAY-TXP-H-S
Fan Tray 1       REV 02   760-024497                      FANTRAY-TXP-H-S
Fan Tray 2       REV 05   760-024502                      FANTRAY-TXP-V-S
Fan Tray 3
Fan Tray 4       REV 05   760-024502                      FANTRAY-TXP-V-S
Fan Tray 5       REV 02   760-024502                      FANTRAY-TXP-V-S

lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 03   710-017247                      CHAS-BP-T1600-S
FPM Display      REV 01   710-021387                      CRAFT-T1600-S
CIP              REV 06   710-002895                      CIP-L-T640-S
PEM 1            Rev 02   740-023211   IPUPAC8KTA         PWR-T1600-4-60-DC-S
SCG 0            REV 15   710-003423                      SCG-T-S
SCG 1            REV 15   710-003423                      SCG-T-S
Routing Engine 0 REV 01   740-026941                      RE-DUO-C1800-8G-S
Routing Engine 1 REV 01   740-026941                      RE-DUO-C1800-8G-S
CB 0             REV 06   710-022597                      CB-LCC-S
CB 1             REV 06   710-022597                      CB-LCC-S
FPC 1            REV 07   710-013035                      T640-FPC3-ES
  PIC 0          REV 05   750-015217                      PC-8GE-TYPE3-SFP-IQ2
  PIC 1          REV 03   750-004424                      PC-1XGE-LR
  PIC 2          REV 01   750-003336                      PC-4OC48-SON-SMSR
FPC 3            REV 12   710-013037                      T1600-FPC4-ES
  PIC 0          REV 02   750-010850                      PD-1OC768-SON-SR
FPC 4            REV 05   710-021534                      T640-FPC1-ES
  PIC 0          REV 04   750-014627                      PB-4OC3-1OC12-SON-SFP
  PIC 1          REV 22   750-005634                      PB-1CHOC12SMIR-QPP
  PIC 2          REV 09   750-002911                      PB-4FE-TX
  PIC 3          REV 08   750-021652                      PB-1CHOC12-STM4-IQE-SFP
FPC 5            REV 07   710-007529                      T640-FPC3
  PIC 0          REV 14   750-009567                      PC-1XGE-XENPAK
  PIC 1          REV 16   750-007141                      PC-10GE-SFP
  PIC 2          REV 12   750-009567                      PC-1XGE-XENPAK
FPC 6            REV 07   710-013035                      T640-FPC3-ES
  PIC 0          REV 09   750-009567                      PC-1XGE-XENPAK
```



```
  PIC 1          REV 06   750-015217                        PC-8GE-TYPE3-SFP-IQ2
  PIC 2          REV 06   750-015217                        PC-8GE-TYPE3-SFP-IQ2
FPC 7            REV 03   710-021540                        T640-FPC2-ES
  PIC 0          REV 13   750-001901                        PB-4OC12-SON-SMIR
  PIC 1          REV 05   750-001900                        PB-1OC48-SON-SMSR
  PIC 2          REV 10   750-008155                        PB-2GE-SFP-QPP
  PIC 3          REV 03   750-014638                        PB-1OC48-SON-B-SFP
SIB 0            REV 07   710-022594                        SIB-TXP-T1600-S
SIB 1            REV 07   710-022594                        SIB-TXP-T1600-S
SIB 3            REV 06   710-022594                        SIB-TXP-T1600-S
SIB 4            REV 08   710-022594                        SIB-TXP-T1600-S
Fan Tray 0                                                  FANTRAY-T-S
Fan Tray 1                                                  FANTRAY-T-S
Fan Tray 2                                                  FANTRAY-TXP-R-S

lcc1-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 04   710-017247                     CHAS-BP-T1600-S
FPM Display      REV 01   710-021387                     CRAFT-T1600-S
CIP              REV 06   710-002895                     CIP-L-T640-S
PEM 0            Rev 02   740-023211   IPUPAC8KTA        PWR-T1600-4-60-DC-S
SCG 0            REV 15   710-003423                     SCG-T-S
SCG 1            REV 15   710-003423                     SCG-T-S
Routing Engine 0 REV 01   740-026941                     RE-DUO-C1800-8G-S
Routing Engine 1 REV 01   740-026941                     RE-DUO-C1800-8G-S
CB 0             REV 06   710-022597                     CB-LCC-S
CB 1             REV 06   710-022597                     CB-LCC-S
FPC 0            REV 02   710-010845                     T640-FPC4-ES
  PIC 0          REV 11   750-017405                     PD-4XGE-XFP
FPC 1            REV 16   710-013037                     T1600-FPC4-ES
  PIC 1          REV 06   750-034781                     PD-1CE-CFP
FPC 2            REV 16   710-013037                     T1600-FPC4-ES
  PIC 1          REV 05   750-034781                     PD-1CE-CFP
FPC 3            REV 10   710-021534                     T640-FPC1-ES
  PIC 0          REV 13   750-012266                     PB-4GE-TYPE1-SFP-IQ2
  PIC 1          REV 01   750-007641                     PE-1GE-SFP-QPP
  PIC 3          REV 17   750-007444                     PB-1CHSTM1-SMIR-QPP
FPC 4            REV 06   710-013035                     T640-FPC3-ES
  PIC 0          REV 22   750-007141                     PC-10GE-SFP
  PIC 1          REV 16   750-009450                     PC-1OC192-SON-SR2
  PIC 2          REV 05   750-004424                     PC-1XGE-LR
  PIC 3          REV 12   750-013423                     PC-MS-500-3
FPC 5            REV 07   710-013560                     T640-FPC3-E2
  PIC 0          REV 11   750-012793                     PC-1XGE-TYPE3-XFP-IQ2
  PIC 1          REV 01   750-004695                     PC-TUNNEL
  PIC 2          REV 32   750-003700                     PC-1OC192-SON-VSR
  PIC 3          REV 12   750-009553                     PC-4OC48-SON-SFP
FPC 6            REV 07   710-013035                     T640-FPC3-ES
  PIC 0          REV 07   750-015217                     PC-8GE-TYPE3-SFP-IQ2
  PIC 1          REV 03   750-003336                     PC-4OC48-SON-SMSR
  PIC 3          REV 02   750-012793                     PC-1XGE-TYPE3-XFP-IQ2
FPC 7            REV 08   710-010845                     T640-FPC4-ES
  PIC 0          REV 11   750-017405                     PD-4XGE-XFP
SIB 0            REV 07   710-022594                     SIB-TXP-T1600-S
SIB 1            REV 07   710-022594                     SIB-TXP-T1600-S
SIB 3            REV 07   710-022594                     SIB-TXP-T1600-S
SIB 4            REV 08   710-022594                     SIB-TXP-T1600-S
Fan Tray 0                                               FANTRAY-T-S
```



```
Fan Tray 1                                            FANTRAY-T-S
Fan Tray 2                                            FANTRAY-TXP-R-S
```

show chassis hardware detail (TX Matrix Plus Router)

```
user@host> show chassis hardware detail
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN111B023AHB      TXP
Midplane         REV 01   710-022574   TR7990            SFC Midplane
FPM Display      REV 03   710-024027   DW4699            TXP FPM Display
CIP 0            REV 01   710-023792   DR1437            TXP CIP
CIP 1            REV 02   710-023792   DS4564            TXP CIP
PEM 0            Rev 07   740-027463   UM26360           Power Entry Module
Routing Engine 0 REV 01   740-026942   737A-1024         SFC RE
  ad0    3887 MB  SMART CF             200811050193CEB1CEB1 Compact Flash
  ad1   30533 MB  SAMSUNG MCBQE32G8MPP-0V SY814A0762     Disk 1
Routing Engine 1 REV 01   740-026942   737A-1024         SFC RE
  ad0    3887 MB  SMART CF             20081105004C19A019A0 Compact Flash
  ad1   30533 MB  SAMSUNG MCBQE32G8MPP-0V SY814A0794     Disk 1
CB 0             REV 03   710-022606   DR7134            SFC Control Board
CB 1             REV 01   710-022606   DP8890            SFC Control Board
SPMB 0                    BUILTIN                        SFC Switch CPU
SPMB 1                    BUILTIN                        SFC Switch CPU
SIB F13 0        REV 03   750-024564   DT9478            F13 SIB
  B Board        REV 02   710-023431   DT6554            F13 SIB
SIB F13 1        REV 03   750-024564   DT9454            F13 SIB
  B Board        REV 02   710-023431   DT6551            F13 SIB
SIB F2S 0/0      REV 02   710-022603   DT2838            F2S SIB
  B Board        REV 02   710-023787   DT1725            NEO PMB
SIB F2S 0/2      REV 02   710-022603   DT2824            F2S SIB
  B Board        REV 02   710-023787   DT1706            NEO PMB
SIB F2S 0/4      REV 02   710-022603   DT2822            F2S SIB
  B Board        REV 02   710-023787   DT1696            NEO PMB
SIB F2S 0/6      REV 02   710-022603   DT2823            F2S SIB
  B Board        REV 02   710-023787   DT1717            NEO PMB
SIB F2S 1/0      REV 03   710-022603   DV0059            F2S SIB
  B Board        REV 03   710-023787   DT9942            NEO PMB
SIB F2S 1/2      REV 02   710-022603   DT2826            F2S SIB
  B Board        REV 02   710-023787   DT1713            NEO PMB
SIB F2S 1/4      REV 03   710-022603   DV0092            F2S SIB
  B Board        REV 03   710-023787   DV0000            NEO PMB
SIB F2S 1/6      REV 03   710-022603   DV0079            F2S SIB
  B Board        REV 03   710-023787   DT9972            NEO PMB
SIB F2S 2/0      REV 03   710-022603   DV0100            F2S SIB
  B Board        REV 03   710-023787   DT9925            NEO PMB
SIB F2S 2/2      REV 03   710-022603   DV0050            F2S SIB
  B Board        REV 03   710-023787   DV0005            NEO PMB
SIB F2S 2/4      REV 03   710-022603   DV0097            F2S SIB
  B Board        REV 03   710-023787   DT9936            NEO PMB
Fan Tray 0       REV 02   760-024497   DR8286            Front Fan Tray
Fan Tray 1       REV 06   760-024497   DV9624            Front Fan Tray
Fan Tray 2       REV 02   760-024502   DR8259            Rear Fan Tray
Fan Tray 3       REV 02   760-024502   DR8270            Rear Fan Tray
Fan Tray 4       REV 02   760-024502   DR8284            Rear Fan Tray
Fan Tray 5       REV 06   760-024502   DV7813            Rear Fan Tray

lcc0-re0:
--------------------------------------------------------------------------
```



```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1101F27AHA      T1600
Midplane         REV 04   710-017247   RC5317            T Series Backplane
FPM GBUS         REV 10   710-002901   DS8197            T640 FPM Board
FPM Display      REV 01   710-021387   DS6433            T1600 FPM Display
CIP              REV 06   710-002895   DS1493            T Series CIP
PEM 0            Rev 08   740-017906   UD26601           Power Entry Module 3x80
SCG 0            REV 15   710-003423   DP5847            T640 Sonet Clock Gen.
SCG 1            REV 15   710-003423   DR0924            T640 Sonet Clock Gen.
Routing Engine 0 REV 01   740-026942   737F-1024         LCC RE
  ad0     3887 MB  SMART CF            2008110502B63E513E51 Compact Flash
  ad1    30533 MB  SAMSUNG MCBQE32G8MPP-0V SY814A1208    Disk 1
Routing Engine 1 REV 01   740-026942   737F-1024         LCC RE
  ad0     3887 MB  SMART CF            2008110500F9A8A8A8A8 Compact Flash
  ad1    30533 MB  SAMSUNG MCBQE32G8MPP-0V SY814A1076    Disk 1
CB 0             REV 05   710-022597   DV4264            LCC Control Board
CB 1             REV 03   710-022597   DP8558            LCC Control Board
FPC 0            REV 14   710-013037   DS9967            FPC Type 4-ES
  CPU            REV 08   710-016744   DS3989            ST-PMB2
  PIC 0          REV 12   750-013198   DL7506            1x Tunnel
  PIC 1          REV 12   750-013198   DL7505            1x Tunnel
  MMB 0          REV 01   710-025563   DS8524            ST-MMB2
  MMB 1          REV 01   710-025563   DS8373            ST-MMB2
FPC 1            REV 14   710-013037   DT0027            FPC Type 4-ES
  CPU            REV 09   710-016744   DS7684            ST-PMB2
  PIC 0          REV 12   750-013198   DL7512            1x Tunnel
  PIC 1          REV 12   750-013198   DL7498            1x Tunnel
  MMB 0          REV 01   710-025563   DS8494            ST-MMB2
  MMB 1          REV 01   710-025563   DS8436            ST-MMB2
SPMB 0           REV 04   710-023321   DV3867            LCC Switch CPU
SPMB 1           REV 02   710-023321   DP0238            LCC Switch CPU
SIB 0            REV 06   710-022594   DT8268            LCC SIB
  B Board        REV 06   710-023185   DT5791            LCC SIB Mezz
SIB 1            REV 06   710-022594   DT8261            LCC SIB
  B Board        REV 06   710-023185   DT5769            LCC SIB Mezz
SIB 2            REV 04   710-022594   DS2315            LCC SIB
  B Board        REV 06   710-023185   DT5788            LCC SIB Mezz
SIB 3            REV 06   710-022594   DT8253            LCC SIB
  B Board        REV 06   710-023185   DT5811            LCC SIB Mezz
SIB 4            REV 06   710-022594   DT8248            LCC SIB
  B Board        REV 06   710-023185   DT5812            LCC SIB Mezz
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Fan Tray
```

show chassis hardware models (TX Matrix Plus Router)

```
user@host> show chassis hardware models
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
FPM Display      REV 03   710-024027   DX0282            CRAFT-TXP
CIP 0            REV 04   710-023792   DW4889            CIP-TXP
CIP 1            REV 04   710-023792   DW4887            CIP-TXP
PEM 0            Rev 07   740-027463   UM26368           ÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿ
Routing Engine 0 REV 01   740-026942   737A-1064         RE-TXP-SFC-DUO-2600-16G
Routing Engine 1 REV 01   740-026942   737A-1082         RE-TXP-SFC-DUO-2600-16G
CB 0             REV 09   710-022606   DW6099            CB-TXP
CB 1             REV 09   710-022606   DW6096            CB-TXP
```



```
SIB F13 1        REV 04   750-024564   DW5776           SIB-TXP-F13
SIB F13 3        REV 04   750-024564   DW5762           SIB-TXP-F13
SIB F13 4        REV 04   750-024564   DW5797           SIB-TXP-F13
SIB F13 6        REV 04   750-024564   DW5770           SIB-TXP-F13
SIB F13 7        REV 04   750-024564   DW5758           SIB-TXP-F13
SIB F13 8        REV 04   750-024564   DW5761           SIB-TXP-F13
SIB F13 9        REV 04   750-024564   DW5754           SIB-TXP-F13
SIB F13 12       REV 04   750-024564   DW5794           SIB-TXP-F13
SIB F2S 0/0      REV 05   710-022603   DW7897
SIB F2S 0/2      REV 05   710-022603   DW7833
SIB F2S 0/4      REV 05   710-022603   DW7875
SIB F2S 0/6      REV 05   710-022603   DW7860
SIB F2S 1/0      REV 04   710-022603   DW4820
SIB F2S 1/2      REV 05   710-022603   DW7849
SIB F2S 1/4      REV 05   710-022603   DW7927           SIB-TXP-F2S
SIB F2S 1/6      REV 05   710-022603   DW7866
SIB F2S 2/0      REV 05   710-022603   DW7880
SIB F2S 2/2      REV 05   710-022603   DW7895
SIB F2S 2/4      REV 05   710-022603   DW7907
SIB F2S 2/6      REV 05   710-022603   DW7785
SIB F2S 3/0      REV 05   710-022603   DW7782
SIB F2S 3/2      REV 05   710-022603   DW7793
SIB F2S 3/4      REV 05   710-022603   DW7779
SIB F2S 3/6      REV 05   710-022603   DW7930
SIB F2S 4/0      REV 05   710-022603   DW7867
SIB F2S 4/2      REV 05   710-022603   DW7917
SIB F2S 4/4      REV 05   710-022603   DW7929
SIB F2S 4/6      REV 05   710-022603   DW7870
Fan Tray 0       REV 06   760-024497   DV7831           FANTRAY-TXP-F
Fan Tray 1       REV 06   760-024497   DV9614           FANTRAY-TXP-F
Fan Tray 2       REV 06   760-024502   DV9618           FANTRAY-TXP-R
Fan Tray 3       REV 06   760-024502   DV9616           FANTRAY-TXP-R
Fan Tray 4       REV 06   760-024502   DV7807           FANTRAY-TXP-R
Fan Tray 5       REV 06   760-024502   DV7828           FANTRAY-TXP-R

lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    FRU model number
Midplane         REV 03   710-017247   RC3765           CHAS-BP-T1600-S
FPM Display      REV 01   710-021387   DN5441           CRAFT-T1600-S
CIP              REV 06   710-002895   DP6021           CIP-L-T640-S
PEM 0            Rev 07   740-017906   UA26384          PWR-T1600-3-80-DC-S
PEM 1            Rev 07   740-017906   UA26296          PWR-T1600-3-80-DC-S
SCG 0            REV 15   710-003423   DR0875           SCG-T-S
CB 0             REV 06   710-022597   DW8534           CB-LCC
CB 1             REV 06   710-022597   DW8527           CB-LCC
FPC 4            REV 12   710-013037   DJ8717           T1600-FPC4-ES
  PIC 0          REV 11   750-017405   DP8795           PD-4XGE-XFP
  PIC 1          REV 11   750-017405   DP8794           PD-4XGE-XFP
FPC 6            REV 14   710-013037   DS5335           T1600-FPC4-ES
  PIC 0          REV 13   750-017405   DS7634           PD-4XGE-XFP
  PIC 1          REV 13   750-017405   DS7637           PD-4XGE-XFP
FPC 7            REV 07   710-013035   DM0990           T1600-FPC3-ES
  PIC 0          REV 16   750-007141   JJ8067           PC-10GE-SFP
  PIC 1          REV 08   750-015749   WE9598           PC-1OC192-SON-XFP
  PIC 2          REV 10   750-009450   HX6466           PC-1OC192-SON-SR2
SIB 0            REV 08   710-022594   DW8033           SIB-TXP-T1600-S
SIB 1            REV 08   710-022594   DW8044           SIB-TXP-T1600-S
SIB 2            REV 08   710-022594   DW8020           SIB-TXP-T1600-S
SIB 3            REV 08   710-022594   DW8063           SIB-TXP-T1600-S
```



```
SIB 4            REV 08   710-022594   DW8064           SIB-TXP-T1600-S
Fan Tray 0                                              FANTRAY-T-S
Fan Tray 1                                              FANTRAY-T-S
Fan Tray 2                                              FANTRAY-TXP-R-S

lcc1-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    FRU model number
Midplane         REV 04   710-017247   RC5361           CHAS-BP-T1600-S
FPM Display      REV 01   710-021387   DS6430           CRAFT-T1600-S
CIP              REV 06   710-002895   DS4239           CIP-L-T640-S
PEM 0            Rev 08   740-017906   UD26649          PWR-T1600-3-80-DC-S
SCG 0            REV 15   710-003423   DP5820           SCG-T-S
CB 0             REV 06   710-022597   DW8523           CB-LCC
CB 1             REV 06   710-022597   DW8528           CB-LCC
FPC 4            REV 12   710-013037   DP8509           T1600-FPC4-ES
  PIC 0          REV 11   750-017405   DP8808           PD-4XGE-XFP
  PIC 1          REV 11   750-017405   DP7263           PD-4XGE-XFP
FPC 6            REV 14   710-013037   DS9961           T1600-FPC4-ES
  PIC 0          REV 13   750-017405   DS5532           PD-4XGE-XFP
  PIC 1          REV 13   750-017405   DS7639           PD-4XGE-XFP
FPC 7            REV 03   710-013035   DF5564           T1600-FPC3-ES
  PIC 0          REV 16   750-007141   JJ8063           PC-10GE-SFP
SIB 0            REV 08   710-022594   DW8035           SIB-TXP-T1600-S
SIB 1            REV 10   710-022594   DX7672           SIB-TXP-T1600-S
SIB 2            REV 08   710-022594   DW8060           SIB-TXP-T1600-S
SIB 3            REV 08   710-022594   DW8072           SIB-TXP-T1600-S
SIB 4            REV 08   710-022594   DW8043           SIB-TXP-T1600-S
Fan Tray 0                                              FANTRAY-T-S
Fan Tray 1                                              FANTRAY-T-S
Fan Tray 2                                              FANTRAY-TXP-R-S

lcc2-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    FRU model number
Midplane         REV 03   710-017247   RC3956           CHAS-BP-T1600-S
FPM Display      REV 01   710-021387   DN7030           CRAFT-T1600-S
CIP              REV 06   710-002895   DM3962           CIP-L-T640-S
PEM 0            Rev 08   740-017906   UD26519          PWR-T1600-3-80-DC-S
PEM 1            Rev 07   740-017906   UC26601          PWR-T1600-3-80-DC-S
SCG 0            REV 15   710-003423   DP0277           SCG-T-S
CB 0             REV 06   710-022597   DW8524           CB-LCC
CB 1             REV 06   710-022597   DW8536           CB-LCC
FPC 4            REV 12   710-013037   DR1194           T1600-FPC4-ES
  PIC 0          REV 11   750-017405   DP8811           PD-4XGE-XFP
  PIC 1          REV 11   750-017405   DP8823           PD-4XGE-XFP
FPC 5            REV 12   710-013037   DR1184           T1600-FPC4-ES
  PIC 1          REV 11   750-017405   DP4744           PD-4XGE-XFP
FPC 6            REV 12   710-013037   DN8622           T1600-FPC4-ES
  PIC 0          REV 14   750-012518   JY9924           PD-4OC192-SON-XFP
  PIC 1          REV 11   750-017405   DP8776           PD-4XGE-XFP
FPC 7            REV 04   710-013560   JR3968           T640-FPC3-E2
  PIC 0          REV 16   750-007141   NC9330           PC-10GE-SFP
SIB 0            REV 07   710-022594   DW4217           SIB-TXP-T1600-S
SIB 1            REV 07   710-022594   DW4213           SIB-TXP-T1600-S
SIB 2            REV 07   710-022594   DW4189           SIB-TXP-T1600-S
SIB 3            REV 07   710-022594   DW4173           SIB-TXP-T1600-S
SIB 4            REV 07   710-022594   DW4201           SIB-TXP-T1600-S
Fan Tray 0                                              FANTRAY-T-S
```

```
Fan Tray 1                                                FANTRAY-T-S
Fan Tray 2                                                FANTRAY-TXP-R-S

lcc3-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number     FRU model number
Midplane          REV 04   710-017247   RC5319            CHAS-BP-T1600-S
FPM Display       REV 01   710-021387   DS6402            CRAFT-T1600-S
CIP               REV 06   710-002895   DR9973            CIP-L-T640-S
PEM 0             Rev 07   740-017906   UC26496           PWR-T1600-3-80-DC-S
PEM 1             Rev 07   740-017906   UC26599           PWR-T1600-3-80-DC-S
SCG 0             REV 15   710-003423   DP5831            SCG-T-S
CB 0              REV 06   710-022597   DW8533            CB-LCC
CB 1              REV 06   710-022597   DW8538            CB-LCC
FPC 0             REV 14   710-013037   DS5345            T1600-FPC4-ES
  PIC 0           REV 13   750-017405   DS7641            PD-4XGE-XFP
  PIC 1           REV 13   750-017405   DS5479            PD-4XGE-XFP
FPC 1             REV 14   710-013037   DS7338            T1600-FPC4-ES
  PIC 0           REV 13   750-017405   DS7631            PD-4XGE-XFP
  PIC 1           REV 13   750-017405   DS7632            PD-4XGE-XFP
FPC 2             REV 14   710-013037   DS9962            T1600-FPC4-ES
  PIC 0           REV 13   750-017405   DS7581            PD-4XGE-XFP
  PIC 1           REV 13   750-017405   DS7627            PD-4XGE-XFP
FPC 4             REV 10   710-010845   JZ6573            T640-FPC4-ES
  PIC 0           REV 14   750-012518   JT5124            PD-4OC192-SON-XFP
FPC 5             REV 14   710-013037   DT0016            T1600-FPC4-ES
  PIC 0           REV 14   750-012518   JY9918            PD-4OC192-SON-XFP
FPC 7             REV 07   710-013035   DM0967            T1600-FPC3-ES
  PIC 0           REV 16   750-007141   JJ8059            PC-10GE-SFP
  PIC 1           REV 13   750-004695   DM5712            PC-TUNNEL
SIB 0             REV 07   710-022594   DW4174            SIB-TXP-T1600-S
SIB 1             REV 07   710-022594   DW4207            SIB-TXP-T1600-S
SIB 2             REV 06   710-022594   DT8231            SIB-TXP-T1600-S
SIB 3             REV 07   710-022594   DW4175            SIB-TXP-T1600-S
SIB 4             REV 07   710-022594   DW4209            SIB-TXP-T1600-S
Fan Tray 0                                                FANTRAY-T-S
Fan Tray 1                                                FANTRAY-T-S
Fan Tray 2                                                FANTRAY-TXP-R-S
```

show chassis hardware (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item              Version  Part number  Serial number     Description
Chassis                                 JN11CAAA4AHB      TXP
Midplane          REV 05   710-022574   ABAC4696          SFC Midplane
FPM Display       REV 09   710-024027   EH3138            TXP FPM Display
CIP 0             REV 12   710-023792   EF6349            TXP CIP
CIP 1             REV 12   710-023792   EG5294            TXP CIP
PEM 0             Rev 06   740-027463   XH04595           Power Entry Module
PEM 1             Rev 06   740-027463   XH04592           Power Entry Module
Routing Engine 0  REV 07   740-026942   P737A-002541      RE-DUO-2600
Routing Engine 1  REV 07   740-026942   P737A-002602      RE-DUO-2600
CB 0              REV 15   710-022606   EH4376            SFC Control Board
CB 1              REV 15   710-022606   EH4379            SFC Control Board
SPMB 0                     BUILTIN                        SFC Switch CPU
SPMB 1                     BUILTIN                        SFC Switch CPU
SIB F13 0         REV 10   750-035002   EM9305            F13 SIB 3D
```



```
  B Board        REV 06   711-035082   EM9667           F13 SIB 3D Mezz
  P Board        REV 05   711-043544   EM9708           F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB34FB00S        CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01H        CXP Module
  Xcvr 4         REV 01   740-047547   XB34FB02W        CXP Module
  Xcvr 6         REV 01   740-047547   XB34FB01T        CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB00W        CXP Module
  Xcvr 10        REV 01   740-047547   XB34FB01S        CXP Module
  Xcvr 12        REV 01   740-047547   XB34FB03H        CXP Module
  Xcvr 14        REV 01   740-047547   XB34FB023        CXP Module
SIB F13 3        REV 01   710-035001   EJ2612           F13 SIB 3D
  B Board        REV 01   711-035082   EJ3815           F13 SIB 3D Mezz
  P Board        REV 01   711-043544   EJ2678           F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB48FB04C        CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB00Z        CXP Module
  Xcvr 4         REV 01   740-047547   XB47FB036        CXP Module
  Xcvr 6         REV 01   740-047547   XB47FB029        CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB02N        CXP Module
  Xcvr 10        REV 01   740-047547   XB42FB0CS        CXP Module
  Xcvr 12        REV 01   740-047547   XB47FB01X        CXP Module
  Xcvr 14        REV 01   740-047547   XB48FB02F        CXP Module
SIB F13 6        REV 05   750-035002   EK2675           F13 SIB 3D
  B Board        REV 03   711-035082   EK2612           F13 SIB 3D Mezz
  P Board        REV 04   711-043544   EK1179           F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB48FB01T        CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB02M        CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB031        CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB04P        CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB02T        CXP Module
  Xcvr 10        REV 01   740-047547   XB34FB01V        CXP Module
  Xcvr 12        REV 01   740-047547   XB48FB02C        CXP Module
  Xcvr 14                 NON-JNPR                      No Module
SIB F13 12       REV 01   710-035001   EJ2631           F13 SIB 3D
  B Board        REV 01   711-035082   EJ3808           F13 SIB 3D Mezz
  P Board        REV 01   711-043544   EJ2676           F13 SIB 3D Power
SIB F2S 0/0      REV 01   711-034977   EH9829           F2S SIB 3D
  B Board        REV 01   711-034979   EH9927           F2S SIB 3D Mezz
SIB F2S 0/2      REV 01   711-034977   EH9791           F2S SIB 3D
  B Board        REV 01   711-034979   EH9852           F2S SIB 3D Mezz
SIB F2S 0/4      REV 01   711-034977   EH9803           F2S SIB 3D
  B Board        REV 01   711-034979   EH9915           F2S SIB 3D Mezz
SIB F2S 0/6      REV 01   711-034977   EH9763           F2S SIB 3D
  B Board        REV 01   711-034979   EH9880           F2S SIB 3D Mezz
SIB F2S 1/0      REV 01   711-034977   EH9757           F2S SIB 3D
  B Board        REV 01   711-034979   EH9889           F2S SIB 3D Mezz
SIB F2S 1/2      REV 01   711-034977   EH9815           F2S SIB 3D
  B Board        REV 01   711-034979   EH9890           F2S SIB 3D Mezz
SIB F2S 1/4      REV 08   750-034978   EN1954           F2S SIB 3D
  B Board        REV 02   711-034979   EN1436           F2S SIB 3D Mezz
SIB F2S 1/6      REV 01   711-034977   EJ7054           F2S SIB 3D
  B Board        REV 01   711-034979   EJ8238           F2S SIB 3D Mezz
SIB F2S 2/0      REV 01   711-034977   EH9830           F2S SIB 3D
  B Board        REV 01   711-034979   EH9844           F2S SIB 3D Mezz
SIB F2S 2/2      REV 01   711-034977   EH9818           F2S SIB 3D
  B Board        REV 01   711-034979   EH9888           F2S SIB 3D Mezz
SIB F2S 2/4      REV 01   711-034977   EH9795           F2S SIB 3D
  B Board        REV 01   711-034979   EH9869           F2S SIB 3D Mezz
SIB F2S 2/6      REV 01   711-034977   EJ7026           F2S SIB 3D
  B Board        REV 01   711-034979   EJ8273           F2S SIB 3D Mezz
SIB F2S 3/0      REV 01   711-034977   EH9811           F2S SIB 3D
  B Board        REV 01   711-034979   EH9892           F2S SIB 3D Mezz
```

```
SIB F2S 3/2      REV 01   711-034977   EH9812           F2S SIB 3D
  B Board        REV 01   711-034979   EH9877           F2S SIB 3D Mezz
SIB F2S 3/4      REV 08   750-034978   EN1947           F2S SIB 3D
  B Board        REV 02   711-034979   EN1471           F2S SIB 3D Mezz
Fan Tray 0       REV 10   760-024497   EH3313           Front Fan Tray
Fan Tray 1       REV 10   760-024497   EH3290           Front Fan Tray
Fan Tray 2       REV 10   760-024502   EH3292           Rear Fan Tray
Fan Tray 3       REV 10   760-024502   EH3287           Rear Fan Tray
Fan Tray 4       REV 10   760-024502   EH3286           Rear Fan Tray
Fan Tray 5       REV 10   760-024502   EH3285           Rear Fan Tray

lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN11B23FEAHA     T1600
Midplane         REV 01   710-027486   RC9787           T-series Backplane
FPM GBUS         REV 13   710-002901   BBAG5132         T640 FPM Board
FPM Display      REV 04   710-021387   BBAL9612         T1600 FPM Display
CIP              REV 06   710-002895   BBAN0605         T-series CIP
PEM 0            REV 05   740-036442   1G022060143      Power Entry Module 6x60
PEM 1            REV 05   740-036442   1G022060011      Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAL7318         T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAL7255         T640 Sonet Clock Gen.
Routing Engine 0 REV 07   740-026941   P737F-002933     RE-DUO-1800
Routing Engine 1 REV 06   740-026941   P737F-002749     RE-DUO-1800
CB 0             REV 11   710-022597   EH3611           LCC Control Board
CB 1             REV 11   710-022597   EH4798           LCC Control Board
FPC 5            REV 17   710-013037   BBAC5333         FPC Type 4-ES
  CPU            REV 10   710-016744   BBAB7619         ST-PMB2
  PIC 0          REV 18   750-017405   BBAE3420         4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 03   740-014289   T10C90659        XFP-10G-SR
  MMB 0          REV 05   710-025563   BBAB9538         ST-MMB2
  MMB 1          REV 05   710-025563   BBAB9502         ST-MMB2
FPC 7            REV 01   750-045173   BBAV0032         FPC Type 5-3D
  CPU
SPMB 0           REV 05   710-023321   EG9434           LCC Switch CPU
SPMB 1           REV 05   710-023321   EH3878           LCC Switch CPU
SIB 0            REV 01   750-041657   EH7997           LCC SIB 3D
  B Board        REV 01   711-042424   EH7674           LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB014        CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB05A        CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB052        CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB01B        CXP Module
SIB 1            REV 01   750-041657   EH8023           LCC SIB 3D
  B Board        REV 01   711-042424   EH7659           LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05J        CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01E        CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB01J        CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB02S        CXP Module
SIB 2            REV 03   750-041657   EJ6554           LCC SIB 3D
  B Board        REV 02   711-042424   EJ5756           LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB34FB01Z        CXP Module
  Xcvr 2         REV 01   740-047547   XB34FB013        CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB04Z        CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB05N        CXP Module
Fan Tray 0                                              Front Top Fan Tray
Fan Tray 1                                              Front Bottom Fan Tray
Fan Tray 2                                              Rear Fan Tray -- Rev 4

lcc2-re0:
```



```
---------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11B3975AHA      T1600
Midplane         REV 01   710-027486   RC9826            T-series Backplane
FPM GBUS         REV 13   710-002901   BBAG5124          T640 FPM Board
FPM Display      REV 03   710-021387   BBAJ1112          T1600 FPM Display
CIP              REV 06   710-002895   BBAL3744          T-series CIP
PEM 0            REV 05   740-036442   1G022060081       Power Entry Module 6x60
PEM 1            REV 05   740-036442   1G022060188       Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAH8775          T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAL7272          T640 Sonet Clock Gen.
Routing Engine 0 REV 07   740-026941   P737F-002992      RE-DUO-1800
Routing Engine 1 REV 07   740-026941   P737F-002938      RE-DUO-1800
CB 0             REV 11   710-022597   EH4805            LCC Control Board
CB 1             REV 11   710-022597   EH4786            LCC Control Board
FPC 1            REV 01   710-033873   BBAH0320          FPC Type 3-ES
  CPU            REV 11   710-016744   BBAF3281          ST-PMB2
  MMB 0          REV 06   710-025563   BBAF5061          ST-MMB2
FPC 5            REV 04   710-033871   BBAM5070          FPC Type 4-ES
  CPU            REV 11   710-016744   BBAM6653          ST-PMB2
  PIC 1          REV 20   750-017405   BBAM1296          4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 03   740-014289   T10B42981         XFP-10G-SR
  MMB 0          REV 07   710-025563   BBAN2631          ST-MMB2
  MMB 1          REV 07   710-025563   BBAN2538          ST-MMB2
SPMB 0           REV 05   710-023321   EH3903            LCC Switch CPU
SPMB 1           REV 05   710-023321   EH3902            LCC Switch CPU
SIB 0            REV 01   750-041657   EH8019            LCC SIB 3D
  B Board        REV 01   711-042424   EH7680            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB04F         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB04S         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB04B         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB043         CXP Module
SIB 1            REV 01   750-041657   EH8012            LCC SIB 3D
  B Board        REV 01   711-042424   EH7658            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05E         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01Z         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB018         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB054         CXP Module
SIB 2            REV 01   750-041657   EH7993            LCC SIB 3D
  B Board        REV 01   711-042424   EH7678            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05C         CXP Module
  Xcvr 2         REV 01   740-047547   XB47FB00N         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB05U         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB05L         CXP Module
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Fan Tray -- Rev 4
```

show chassis hardware clei-models (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware clei-models
sfc0-re0:
---------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 05   710-022574                     CHAS-BP-TXP-S
FPM Display      REV 09   710-024027                     CRAFT-TXP-S
CIP 0            REV 12   710-023792                     CIP-TXP-S
CIP 1            REV 12   710-023792                     CIP-TXP-S
PEM 0            Rev 06   740-027463   IPUPAFGKTA        PWR-TXP-7-60-DC-S
```

```
Routing Engine 0 REV 07   740-026942                    RE-DUO-C2600-16G-S
Routing Engine 1 REV 07   740-026942                    RE-DUO-C2600-16G-S
CB 0             REV 13   710-022606                    CB-TXP-S
CB 1             REV 14   710-022606                    CB-TXP-S
SIB F13 0        REV 10   750-035002   PROTOXCLEI       SIB-TXP-3D-F13-S
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
  Xcvr 8         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 10        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 12        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 14        REV 01   740-047547                    CXP-TXP-3D
SIB F13 1        REV 10   750-035002   PROTOXCLEI       SIB-TXP-3D-F13-S
  Xcvr 0         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 1         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 2         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 3         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 4         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 5         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 6         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 7         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 8         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 10        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 12        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 14        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
  Xcvr 8         REV 01   740-048813
  Xcvr 10        REV 01   740-048813
  Xcvr 12        REV 01   740-048813
  Xcvr 14        REV 01   740-048813
  Xcvr 0         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 1         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 2         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 3         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 4         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 5         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 6         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 7         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 8         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 10        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 12        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 14        REV 01   740-047547                    CXP-TXP-3D
SIB F13 6        REV 16   750-035002   PROTOXCLEI       SIB-TXP-3D-F13
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
```



```
  Xcvr 6          REV 01   740-048813
  Xcvr 7          REV 01   740-048813
  Xcvr 8          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 10         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 12         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 14         REV 01   740-047547                  CXP-TXP-3D
SIB F13 7         REV 10   750-035002   PROTOXCLEI     SIB-TXP-3D-F13-S
  Xcvr 0          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 1          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 2          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 3          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 4          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 5          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 6          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 7          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 8          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 10         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 12         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 14         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 0          REV 01   740-048813
  Xcvr 1          REV 01   740-048813
  Xcvr 2          REV 01   740-048813
  Xcvr 3          REV 01   740-048813
  Xcvr 4          REV 01   740-048813
  Xcvr 5          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 6          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 7          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 8          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 10         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 12         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 14         REV 01   740-047547                  CXP-TXP-3D
SIB F13 9         REV 16   750-035002   PROTOXCLEI     SIB-TXP-3D-F13
  Xcvr 0          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 1          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 2          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 3          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 4          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 5          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 6          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 7          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 8          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 10         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 12         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 14         REV 01   740-047547                  CXP-TXP-3D
SIB F13 11        REV 10   750-035002   PROTOXCLEI     750-035002
  Xcvr 0          REV 01   740-048813
  Xcvr 1          REV 01   740-048813
  Xcvr 2          REV 01   740-048813
  Xcvr 3          REV 01   740-048813
  Xcvr 4          REV 01   740-048813
  Xcvr 5          REV 01   740-048813
  Xcvr 6          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 7          REV 01   740-048813
  Xcvr 8          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 12         REV 01   740-047547                  CXP-TXP-3D
  Xcvr 14         REV 01   740-047547                  CXP-TXP-3D
SIB F13 12        REV 16   750-035002   PROTOXCLEI     SIB-TXP-3D-F13
  Xcvr 0          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 1          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 2          REV 01   740-047547                  CXP-TXP-3D
  Xcvr 3          REV 01   740-047547                  CXP-TXP-3D
```

```
  Xcvr 4         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 5         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 6         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 7         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 8         REV 01   740-047547                    CXP-TXP-3D
  Xcvr 10        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 12        REV 01   740-047547                    CXP-TXP-3D
  Xcvr 14        REV 01   740-047547                    CXP-TXP-3D
SIB F2S 0/0      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 0/2      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 0/4      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 0/6      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 1/0      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 1/2      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 1/4      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 1/6      REV 08   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 2/0      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 2/2      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 2/4      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 2/6      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 3/0      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 3/2      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 3/4      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 3/6      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 4/0      REV 07   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 4/2      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 4/4      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
SIB F2S 4/6      REV 06   750-034978   PROTOXCLEI       SIB-TXP-3D-F2S
Fan Tray 0       REV 10   760-024497                    FANTRAY-TXP-H-S
Fan Tray 1       REV 10   760-024497                    FANTRAY-TXP-H-S
Fan Tray 2       REV 10   760-024502                    FANTRAY-TXP-V-S
Fan Tray 3       REV 10   760-024502                    FANTRAY-TXP-V-S
Fan Tray 4       REV 10   760-024502                    FANTRAY-TXP-V-S
Fan Tray 5       REV 10   760-024502                    FANTRAY-TXP-V-S

lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code        FRU model number
Midplane         REV 01   710-027486   IPMJ700DRD       CHAS-BP-T1600-S
FPM Display      REV 04   710-021387                    CRAFT-T1600-S
CIP              REV 06   710-002895                    CIP-L-T640-S
PEM 0            REV 05   740-036442   IPUPAG6KAA       PWR-T-6-60-DC-S
PEM 1            REV 05   740-036442   IPUPAG6KAA       PWR-T-6-60-DC-S
SCG 0            REV 18   710-003423                    SCG-T-S
SCG 1            REV 18   710-003423                    SCG-T-S
Routing Engine 0 REV 10   740-026941                    RE-DUO-C1800-8G-S
Routing Engine 1 REV 07   740-026941                    RE-DUO-C1800-8G-S
CB 0             REV 11   710-022597                    CB-LCC-S
CB 1             REV 11   710-022597                    CB-LCC-S
FPC 0            REV 01   750-045173   IP9IAL4DAB       T4000-FPC5-3D
  PIC 0          REV 17   750-034624   IP9IAL2DAA       PF-12XGE-SFPP
  PIC 1          REV 17   750-034624   IP9IAL2DAA       PF-12XGE-SFPP
FPC 3            REV 01   750-045173   IP9IAL4DAB       T4000-FPC5-3D
  PIC 0          REV 13   750-033423   XXXXXXXXDD       PF-12-24XGE-SFPP
FPC 4            REV 02   750-045173   IP9IAL4DAC       T4000-FPC5-3D
  PIC 0          REV 17   750-034624   IP9IAL2DAA       PF-12XGE-SFPP
  PIC 1          REV 17   750-034624   IP9IAL2DAA       PF-12XGE-SFPP
FPC 5            REV 01   750-045173   IP9IAL4DAB       T4000-FPC5-3D
  PIC 0          REV 17   750-034624   IP9IAL2DAA       PF-12XGE-SFPP
  PIC 1          REV 17   750-034624   IP9IAL2DAA       PF-12XGE-SFPP
```



```
FPC 6            REV 01   750-045173   IP9IAL4DAB        T4000-FPC5-3D
  PIC 0          REV 17   750-034624   IP9IAL2DAA        PF-12XGE-SFPP
  PIC 1          REV 10   750-035293   IP9IAL3DAA        PF-1CGE-CFP
SIB 0            REV 06   750-041657   PROTOXCLEI        SIB-TXP-3D-LCC
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
SIB 1            REV 06   750-041657   PROTOXCLEI        SIB-TXP-3D-LCC
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
SIB 2            REV 06   750-041657   PROTOXCLEI        SIB-TXP-3D-LCC
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
SIB 3            REV 07   750-041657   PROTOXCLEI        SIB-TXP-3D-LLC
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
SIB 4            REV 06   750-041657   PROTOXCLEI        SIB-TXP-3D-LCC
  Xcvr 0         REV 01   740-048813
  Xcvr 1         REV 01   740-048813
  Xcvr 2         REV 01   740-048813
  Xcvr 3         REV 01   740-048813
  Xcvr 4         REV 01   740-048813
  Xcvr 5         REV 01   740-048813
  Xcvr 6         REV 01   740-048813
  Xcvr 7         REV 01   740-048813
Fan Tray 0                                               FANTRAY-T-S
Fan Tray 1                                               FANTRAY-T-S
Fan Tray 2                                               FANTRAY-TXP3D-LCC-R-S
[Output Truncated]
```

show chassis hardware detail (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware detail
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11CAAA4AHB      TXP
```

```
Midplane         REV 05   710-022574   ABAC4696          SFC Midplane
FPM Display      REV 09   710-024027   EH3138            TXP FPM Display
CIP 0            REV 12   710-023792   EF6349            TXP CIP
CIP 1            REV 12   710-023792   EG5294            TXP CIP
PEM 0            Rev 06   740-027463   XH04595           Power Entry Module
PEM 1            Rev 06   740-027463   XH04592           Power Entry Module
Routing Engine 0 REV 07   740-026942   P737A-002541      RE-DUO-2600
  ad0    3823 MB  SMART CF             2011030400062C132C13 Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 201105100009A452A452 Disk 1
Routing Engine 1 REV 07   740-026942   P737A-002602      RE-DUO-2600
  ad0    3823 MB  SMART CF             20110508085EE471E471 Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 201110210089DF39DF39 Disk 1
CB 0             REV 15   710-022606   EH4376            SFC Control Board
CB 1             REV 15   710-022606   EH4379            SFC Control Board
SPMB 0                    BUILTIN                        SFC Switch CPU
SPMB 1                    BUILTIN                        SFC Switch CPU
SIB F13 0        REV 10   750-035002   EM9305            F13 SIB 3D
  B Board        REV 06   711-035082   EM9667            F13 SIB 3D Mezz
  P Board        REV 05   711-043544   EM9708            F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB34FB00S         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01H         CXP Module
  Xcvr 4         REV 01   740-047547   XB34FB02W         CXP Module
  Xcvr 6         REV 01   740-047547   XB34FB01T         CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB00W         CXP Module
  Xcvr 10        REV 01   740-047547   XB34FB01S         CXP Module
  Xcvr 12        REV 01   740-047547   XB34FB03H         CXP Module
  Xcvr 14        REV 01   740-047547   XB34FB023         CXP Module
SIB F13 3        REV 01   710-035001   EJ2612            F13 SIB 3D
  B Board        REV 01   711-035082   EJ3815            F13 SIB 3D Mezz
  P Board        REV 01   711-043544   EJ2678            F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB48FB04C         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB00Z         CXP Module
  Xcvr 4         REV 01   740-047547   XB47FB036         CXP Module
  Xcvr 6         REV 01   740-047547   XB47FB029         CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB02N         CXP Module
  Xcvr 10        REV 01   740-047547   XB42FB0CS         CXP Module
  Xcvr 12        REV 01   740-047547   XB47FB01X         CXP Module
  Xcvr 14        REV 01   740-047547   XB48FB02F         CXP Module
SIB F13 6        REV 05   750-035002   EK2675            F13 SIB 3D
  B Board        REV 03   711-035082   EK2612            F13 SIB 3D Mezz
  P Board        REV 04   711-043544   EK1179            F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB48FB01T         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB02M         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB031         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB04P         CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB02T         CXP Module
  Xcvr 10        REV 01   740-047547   XB34FB01V         CXP Module
  Xcvr 12        REV 01   740-047547   XB48FB02C         CXP Module
  Xcvr 14                 NON-JNPR                       No Module
SIB F13 12       REV 01   710-035001   EJ2631            F13 SIB 3D
  B Board        REV 01   711-035082   EJ3808            F13 SIB 3D Mezz
  P Board        REV 01   711-043544   EJ2676            F13 SIB 3D Power
SIB F2S 0/0      REV 01   711-034977   EH9829            F2S SIB 3D
  B Board        REV 01   711-034979   EH9927            F2S SIB 3D Mezz
SIB F2S 0/2      REV 01   711-034977   EH9791            F2S SIB 3D
  B Board        REV 01   711-034979   EH9852            F2S SIB 3D Mezz
SIB F2S 0/4      REV 01   711-034977   EH9803            F2S SIB 3D
  B Board        REV 01   711-034979   EH9915            F2S SIB 3D Mezz
SIB F2S 0/6      REV 01   711-034977   EH9763            F2S SIB 3D
  B Board        REV 01   711-034979   EH9880            F2S SIB 3D Mezz
SIB F2S 1/0      REV 01   711-034977   EH9757            F2S SIB 3D
```



```
  B Board        REV 01   711-034979   EH9889           F2S SIB 3D Mezz
SIB F2S 1/2      REV 01   711-034977   EH9815           F2S SIB 3D
  B Board        REV 01   711-034979   EH9890           F2S SIB 3D Mezz
SIB F2S 1/4      REV 08   750-034978   EN1954           F2S SIB 3D
  B Board        REV 02   711-034979   EN1436           F2S SIB 3D Mezz
SIB F2S 1/6      REV 01   711-034977   EJ7054           F2S SIB 3D
  B Board        REV 01   711-034979   EJ8238           F2S SIB 3D Mezz
SIB F2S 2/0      REV 01   711-034977   EH9830           F2S SIB 3D
  B Board        REV 01   711-034979   EH9844           F2S SIB 3D Mezz
SIB F2S 2/2      REV 01   711-034977   EH9818           F2S SIB 3D
  B Board        REV 01   711-034979   EH9888           F2S SIB 3D Mezz
SIB F2S 2/4      REV 01   711-034977   EH9795           F2S SIB 3D
  B Board        REV 01   711-034979   EH9869           F2S SIB 3D Mezz
SIB F2S 2/6      REV 01   711-034977   EJ7026           F2S SIB 3D
  B Board        REV 01   711-034979   EJ8273           F2S SIB 3D Mezz
SIB F2S 3/0      REV 01   711-034977   EH9811           F2S SIB 3D
  B Board        REV 01   711-034979   EH9892           F2S SIB 3D Mezz
SIB F2S 3/2      REV 01   711-034977   EH9812           F2S SIB 3D
  B Board        REV 01   711-034979   EH9877           F2S SIB 3D Mezz
SIB F2S 3/4      REV 08   750-034978   EN1947           F2S SIB 3D
  B Board        REV 02   711-034979   EN1471           F2S SIB 3D Mezz
Fan Tray 0       REV 10   760-024497   EH3313           Front Fan Tray
Fan Tray 1       REV 10   760-024497   EH3290           Front Fan Tray
Fan Tray 2       REV 10   760-024502   EH3292           Rear Fan Tray
Fan Tray 3       REV 10   760-024502   EH3287           Rear Fan Tray
Fan Tray 4       REV 10   760-024502   EH3286           Rear Fan Tray
Fan Tray 5       REV 10   760-024502   EH3285           Rear Fan Tray

lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN11B23FEAHA     T1600
Midplane         REV 01   710-027486   RC9787           T-series Backplane
FPM GBUS         REV 13   710-002901   BBAG5132         T640 FPM Board
FPM Display      REV 04   710-021387   BBAL9612         T1600 FPM Display
CIP              REV 06   710-002895   BBAN0605         T-series CIP
PEM 0            REV 05   740-036442   1G022060143      Power Entry Module 6x60
PEM 1            REV 05   740-036442   1G022060011      Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAL7318         T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAL7255         T640 Sonet Clock Gen.
Routing Engine 0 REV 07   740-026941   P737F-002933     RE-DUO-1800
  ad0    3823 MB  SMART CF             201103030490604E604E Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 20110729028B11D411D4 Disk 1
Routing Engine 1 REV 06   740-026941   P737F-002749     RE-DUO-1800
  ad0    3823 MB  SMART CF             2011010504EB99649964 Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 201102140058934A934A Disk 1
CB 0             REV 11   710-022597   EH3611           LCC Control Board
CB 1             REV 11   710-022597   EH4798           LCC Control Board
FPC 5            REV 17   710-013037   BBAC5333         FPC Type 4-ES
  CPU            REV 10   710-016744   BBAB7619         ST-PMB2
  PIC 0          REV 18   750-017405   BBAE3420         4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 03   740-014289   T10C90659        XFP-10G-SR
  MMB 0          REV 05   710-025563   BBAB9538         ST-MMB2
  MMB 1          REV 05   710-025563   BBAB9502         ST-MMB2
FPC 7            REV 01   750-045173   BBAV0032         FPC Type 5-3D
  CPU
SPMB 0           REV 05   710-023321   EG9434           LCC Switch CPU
SPMB 1           REV 05   710-023321   EH3878           LCC Switch CPU
SIB 0            REV 01   750-041657   EH7997           LCC SIB 3D
  B Board        REV 01   711-042424   EH7674           LCC SIB 3D Mezz
```

```
  Xcvr 0         REV 01   740-047547   XB48FB014         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB05A         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB052         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB01B         CXP Module
SIB 1            REV 01   750-041657   EH8023            LCC SIB 3D
  B Board        REV 01   711-042424   EH7659            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05J         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01E         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB01J         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB02S         CXP Module
SIB 2            REV 03   750-041657   EJ6554            LCC SIB 3D
  B Board        REV 02   711-042424   EJ5756            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB34FB01Z         CXP Module
  Xcvr 2         REV 01   740-047547   XB34FB013         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB04Z         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB05N         CXP Module
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Fan Tray -- Rev 4

lcc2-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11B3975AHA      T1600
Midplane         REV 01   710-027486   RC9826            T-series Backplane
FPM GBUS         REV 13   710-002901   BBAG5124          T640 FPM Board
FPM Display      REV 03   710-021387   BBAJ1112          T1600 FPM Display
CIP              REV 06   710-002895   BBAL3744          T-series CIP
PEM 0            REV 05   740-036442   1G022060081       Power Entry Module 6x60
PEM 1            REV 05   740-036442   1G022060188       Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAH8775          T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAL7272          T640 Sonet Clock Gen.
Routing Engine 0 REV 07   740-026941   P737F-002992      RE-DUO-1800
  ad0    3823 MB  SMART CF             201103030356329E329E Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 2011051000488D8B8D8B Disk 1
Routing Engine 1 REV 07   740-026941   P737F-002938      RE-DUO-1800
  ad0    3823 MB  SMART CF             20110304000F02680268 Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 201105300A70F325F325 Disk 1
CB 0             REV 11   710-022597   EH4805            LCC Control Board
CB 1             REV 11   710-022597   EH4786            LCC Control Board
FPC 1            REV 01   710-033873   BBAH0320          FPC Type 3-ES
  CPU            REV 11   710-016744   BBAF3281          ST-PMB2
  MMB 0          REV 06   710-025563   BBAF5061          ST-MMB2
FPC 5            REV 04   710-033871   BBAM5070          FPC Type 4-ES
  CPU            REV 11   710-016744   BBAM6653          ST-PMB2
  PIC 1          REV 20   750-017405   BBAM1296          4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 03   740-014289   T10B42981         XFP-10G-SR
  MMB 0          REV 07   710-025563   BBAN2631          ST-MMB2
  MMB 1          REV 07   710-025563   BBAN2538          ST-MMB2
SPMB 0           REV 05   710-023321   EH3903            LCC Switch CPU
SPMB 1           REV 05   710-023321   EH3902            LCC Switch CPU
SIB 0            REV 01   750-041657   EH8019            LCC SIB 3D
  B Board        REV 01   711-042424   EH7680            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB04F         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB04S         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB04B         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB043         CXP Module
SIB 1            REV 01   750-041657   EH8012            LCC SIB 3D
  B Board        REV 01   711-042424   EH7658            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05E         CXP Module
```



```
  Xcvr 2         REV 01   740-047547   XB48FB01Z         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB018         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB054         CXP Module
SIB 2            REV 01   750-041657   EH7993            LCC SIB 3D
  B Board        REV 01   711-042424   EH7678            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05C         CXP Module
  Xcvr 2         REV 01   740-047547   XB47FB00N         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB05U         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB05L         CXP Module
Fan Tray 0                                               Front Top Fan Tray
Fan Tray 1                                               Front Bottom Fan Tray
Fan Tray 2                                               Rear Fan Tray -- Rev 4
```

show chassis hardware lcc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11B23FEAHA      T1600
Midplane         REV 01   710-027486   RC9787            T-series Backplane
FPM GBUS         REV 13   710-002901   BBAG5132          T640 FPM Board
FPM Display      REV 04   710-021387   BBAL9612          T1600 FPM Display
CIP              REV 06   710-002895   BBAN0605          T-series CIP
PEM 0            REV 05   740-036442   1G022060143       Power Entry Module 6x60
PEM 1            REV 05   740-036442   1G022060011       Power Entry Module 6x60
SCG 0            REV 18   710-003423   BBAL7318          T640 Sonet Clock Gen.
SCG 1            REV 18   710-003423   BBAL7255          T640 Sonet Clock Gen.
Routing Engine 0 REV 07   740-026941   P737F-002933      RE-DUO-1800
Routing Engine 1 REV 06   740-026941   P737F-002749      RE-DUO-1800
CB 0             REV 11   710-022597   EH3611            LCC Control Board
CB 1             REV 11   710-022597   EH4798            LCC Control Board
FPC 5            REV 17   710-013037   BBAC5333          FPC Type 4-ES
  CPU            REV 10   710-016744   BBAB7619          ST-PMB2
  PIC 0          REV 18   750-017405   BBAE3420          4x 10GE (LAN/WAN) XFP
    Xcvr 0       REV 03   740-014289   T10C90659         XFP-10G-SR
  MMB 0          REV 05   710-025563   BBAB9538          ST-MMB2
  MMB 1          REV 05   710-025563   BBAB9502          ST-MMB2
FPC 7            REV 01   750-045173   BBAV0032          FPC Type 5-3D
  CPU
SPMB 0           REV 05   710-023321   EG9434            LCC Switch CPU
SPMB 1           REV 05   710-023321   EH3878            LCC Switch CPU
SIB 0            REV 01   750-041657   EH7997            LCC SIB 3D
  B Board        REV 01   711-042424   EH7674            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB014         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB05A         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB052         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB01B         CXP Module
SIB 1            REV 01   750-041657   EH8023            LCC SIB 3D
  B Board        REV 01   711-042424   EH7659            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB48FB05J         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01E         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB01J         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB02S         CXP Module
SIB 2            REV 03   750-041657   EJ6554            LCC SIB 3D
  B Board        REV 02   711-042424   EJ5756            LCC SIB 3D Mezz
  Xcvr 0         REV 01   740-047547   XB34FB01Z         CXP Module
  Xcvr 2         REV 01   740-047547   XB34FB013         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB04Z         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB05N         CXP Module
```



```
Fan Tray 0                                                Front Top Fan Tray
Fan Tray 1                                                Front Bottom Fan Tray
Fan Tray 2                                                Rear Fan Tray -- Rev 4
```

show chassis hardware sfc (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11CAAA4AHB      TXP
Midplane         REV 05   710-022574   ABAC4696          SFC Midplane
FPM Display      REV 09   710-024027   EH3138            TXP FPM Display
CIP 0            REV 12   710-023792   EF6349            TXP CIP
CIP 1            REV 12   710-023792   EG5294            TXP CIP
PEM 0            Rev 06   740-027463   XH04595           Power Entry Module
PEM 1            Rev 06   740-027463   XH04592           Power Entry Module
Routing Engine 0 REV 07   740-026942   P737A-002541      RE-DUO-2600
Routing Engine 1 REV 07   740-026942   P737A-002602      RE-DUO-2600
CB 0             REV 15   710-022606   EH4376            SFC Control Board
CB 1             REV 15   710-022606   EH4379            SFC Control Board
SPMB 0                    BUILTIN                        SFC Switch CPU
SPMB 1                    BUILTIN                        SFC Switch CPU
SIB F13 0        REV 10   750-035002   EM9305            F13 SIB 3D
  B Board        REV 06   711-035082   EM9667            F13 SIB 3D Mezz
  P Board        REV 05   711-043544   EM9708            F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB34FB00S         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB01H         CXP Module
  Xcvr 4         REV 01   740-047547   XB34FB02W         CXP Module
  Xcvr 6         REV 01   740-047547   XB34FB01T         CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB00W         CXP Module
  Xcvr 10        REV 01   740-047547   XB34FB01S         CXP Module
  Xcvr 12        REV 01   740-047547   XB34FB03H         CXP Module
  Xcvr 14        REV 01   740-047547   XB34FB023         CXP Module
SIB F13 3        REV 01   710-035001   EJ2612            F13 SIB 3D
  B Board        REV 01   711-035082   EJ3815            F13 SIB 3D Mezz
  P Board        REV 01   711-043544   EJ2678            F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB48FB04C         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB00Z         CXP Module
  Xcvr 4         REV 01   740-047547   XB47FB036         CXP Module
  Xcvr 6         REV 01   740-047547   XB47FB029         CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB02N         CXP Module
  Xcvr 10        REV 01   740-047547   XB42FB0CS         CXP Module
  Xcvr 12        REV 01   740-047547   XB47FB01X         CXP Module
  Xcvr 14        REV 01   740-047547   XB48FB02F         CXP Module
SIB F13 6        REV 05   750-035002   EK2675            F13 SIB 3D
  B Board        REV 03   711-035082   EK2612            F13 SIB 3D Mezz
  P Board        REV 04   711-043544   EK1179            F13 SIB 3D Power
  Xcvr 0         REV 01   740-047547   XB48FB01T         CXP Module
  Xcvr 2         REV 01   740-047547   XB48FB02M         CXP Module
  Xcvr 4         REV 01   740-047547   XB48FB031         CXP Module
  Xcvr 6         REV 01   740-047547   XB48FB04P         CXP Module
  Xcvr 8         REV 01   740-047547   XB48FB02T         CXP Module
  Xcvr 10        REV 01   740-047547   XB34FB01V         CXP Module
  Xcvr 12        REV 01   740-047547   XB48FB02C         CXP Module
  Xcvr 14                 NON-JNPR                       No Module
SIB F13 12       REV 01   710-035001   EJ2631            F13 SIB 3D
  B Board        REV 01   711-035082   EJ3808            F13 SIB 3D Mezz
  P Board        REV 01   711-043544   EJ2676            F13 SIB 3D Power
SIB F2S 0/0      REV 01   711-034977   EH9829            F2S SIB 3D
```



```
  B Board        REV 01   711-034979   EH9927            F2S SIB 3D Mezz
SIB F2S 0/2      REV 01   711-034977   EH9791            F2S SIB 3D
  B Board        REV 01   711-034979   EH9852            F2S SIB 3D Mezz
SIB F2S 0/4      REV 01   711-034977   EH9803            F2S SIB 3D
  B Board        REV 01   711-034979   EH9915            F2S SIB 3D Mezz
SIB F2S 0/6      REV 01   711-034977   EH9763            F2S SIB 3D
  B Board        REV 01   711-034979   EH9880            F2S SIB 3D Mezz
SIB F2S 1/0      REV 01   711-034977   EH9757            F2S SIB 3D
  B Board        REV 01   711-034979   EH9889            F2S SIB 3D Mezz
SIB F2S 1/2      REV 01   711-034977   EH9815            F2S SIB 3D
  B Board        REV 01   711-034979   EH9890            F2S SIB 3D Mezz
SIB F2S 1/4      REV 08   750-034978   EN1954            F2S SIB 3D
  B Board        REV 02   711-034979   EN1436            F2S SIB 3D Mezz
SIB F2S 1/6      REV 01   711-034977   EJ7054            F2S SIB 3D
  B Board        REV 01   711-034979   EJ8238            F2S SIB 3D Mezz
SIB F2S 2/0      REV 01   711-034977   EH9830            F2S SIB 3D
  B Board        REV 01   711-034979   EH9844            F2S SIB 3D Mezz
SIB F2S 2/2      REV 01   711-034977   EH9818            F2S SIB 3D
  B Board        REV 01   711-034979   EH9888            F2S SIB 3D Mezz
SIB F2S 2/4      REV 01   711-034977   EH9795            F2S SIB 3D
  B Board        REV 01   711-034979   EH9869            F2S SIB 3D Mezz
SIB F2S 2/6      REV 01   711-034977   EJ7026            F2S SIB 3D
  B Board        REV 01   711-034979   EJ8273            F2S SIB 3D Mezz
SIB F2S 3/0      REV 01   711-034977   EH9811            F2S SIB 3D
  B Board        REV 01   711-034979   EH9892            F2S SIB 3D Mezz
SIB F2S 3/2      REV 01   711-034977   EH9812            F2S SIB 3D
  B Board        REV 01   711-034979   EH9877            F2S SIB 3D Mezz
SIB F2S 3/4      REV 08   750-034978   EN1947            F2S SIB 3D
  B Board        REV 02   711-034979   EN1471            F2S SIB 3D Mezz
Fan Tray 0       REV 10   760-024497   EH3313            Front Fan Tray
Fan Tray 1       REV 10   760-024497   EH3290            Front Fan Tray
Fan Tray 2       REV 10   760-024502   EH3292            Rear Fan Tray
Fan Tray 3       REV 10   760-024502   EH3287            Rear Fan Tray
Fan Tray 4       REV 10   760-024502   EH3286            Rear Fan Tray
Fan Tray 5       REV 10   760-024502   EH3285            Rear Fan Tray
```

show chassis hardware (16-Port 10-Gigabit Ethernet MPC with SFP+ Optics [MX Series Routers])

```
user@host> show chassis hardware
Hardware inventory:
    Item             Version  Part number  Serial number     Description
    Chassis                                JN112D865AFA      MX960
    Midplane         REV 03   710-013698   TS3339            MX960 Backplane
    FPM Board        REV 03   710-014974   WW6267            Front Panel Display

    PDM              Rev 03   740-013110   QCS12485026       Power Distribution
Module
    PEM 0           Rev 04   740-013682   QCS12434086       PS 1.7kW; 200-240VAC
 in
    PEM 1           Rev 04   740-013682   QCS1243408Z       PS 1.7kW; 200-240VAC
 in
    PEM 2           Rev 04   740-013682   QCS1243407X       PS 1.7kW; 200-240VAC
 in
    Routing Engine 0 REV 07   740-015113   9009009677        RE-S-1300
    Routing Engine 1 REV 07   740-015113   9009011510        RE-S-1300
    CB 0             REV 03   710-021523   XF0394            MX SCB
    CB 1             REV 03   710-021523   XF0550            MX SCB
    CB 2             REV 03   710-021523   XD7455            MX SCB
    FPC 4            REV 02   750-028467   JR6127            MPC M 16x 10GE
      CPU            REV 02   711-029089   JX0129            AS PMB
      PIC 0                   BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
```

```
  PIC 1                 BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
  PIC 2                 BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
  PIC 3                 BUILTIN      BUILTIN           4x 10GE(LAN) SFP+
 Fan Tray 0      REV 05  740-014971   TP9990            Fan Tray
 Fan Tray 1      REV 05  740-014971   VS1709            Fan Tray
```

show chassis hardware (MPC3E [MX Series Routers])

```
user@host> show chassis hardware
Hardware inventory:
   Item             Version  Part number   Serial number     Description
   Chassis                                 JN1101AFEAFB      MX480
   Midplane         REV 05   710-017414    TR4444            MX480 Midplane
   FPM Board        REV 02   710-017254    KG6056            Front Panel Display

   PEM 0           Rev 03   740-017330    QCS082090FC      PS 1.2-1.7kW; 100-240V

   PEM 1           Rev 03   740-017330    QCS082090FD      PS 1.2-1.7kW; 100-240V

   Routing Engine 0 REV 07   740-013063    9009004124        RE-S-2000
   Routing Engine 1 REV 07   740-013063    9009005569        RE-S-2000
   CB 0             REV 07   710-021523    XZ3587            MX SCB
   CB 1             REV 03   710-021523    KH8306            MX SCB
   FPC 1            REV 04.1.07 750-033205 P1240             MPC Type 3
    CPU             REV 01   711-035209    YL0504            HMPC PMB 2G
    MIC 1           REV 10   750-033199    YX4495            1X100GE CFP
     PIC 2                   BUILTIN       BUILTIN           1X100GE CFP
      Xcvr 0        REV 01   740-032210    C22CQNE           CFP-100G-LR4
   FPC 2            REV 26   750-016670    KH0045            DPCE 40x 1GE R EQ
    CPU             REV 07   710-013713    KF5448            DPC PMB
    PIC 0                    BUILTIN       BUILTIN           10x 1GE(LAN) EQ
     Xcvr 0         REV 01   740-011613    PF21JHU           SFP-SX
    PIC 1                    BUILTIN       BUILTIN           10x 1GE(LAN) EQ
     Xcvr 9         REV 01   740-011613    AM0813S8ZL6       SFP-SX
    PIC 2                    BUILTIN       BUILTIN           10x 1GE(LAN) EQ
     Xcvr 0         REV 02   740-011613    PGL2KYF           SFP-SX
     Xcvr 2         REV 01   740-011613    AM0806S8N4P       SFP-SX
    PIC 3                    BUILTIN       BUILTIN           10x 1GE(LAN) EQ
     Xcvr 5         REV 01   740-011613    AM0815S967N       SFP-SX
     Xcvr 7         REV 01   740-011613    AM0806S8N1X       SFP-SX
     Xcvr 8         REV 01   740-011613    AM0815S967J       SFP-SX
     Xcvr 9         REV 01   740-011613    AM0815S967M       SFP-SX
   FPC 3            REV 12.2.09 750-033205 YR9443            MPC Type 3
    CPU             REV 03   711-035209    YL6931            HMPC PMB 2G
    MIC 0           REV 05   750-033199    YR3269            1X100GE CFP
     PIC 0                   BUILTIN       BUILTIN           1X100GE CFP
      Xcvr 0        REV 01   740-032210    ULHOKG3           CFP-100G-LR4
    MIC 1           REV 02   750-033199    YG3245            1X100GE CFP
     PIC 2                   BUILTIN       BUILTIN           1X100GE CFP
      Xcvr 0        REV 01   740-032210    ULHOKGF           CFP-100G-LR4
  FPC 4             REV 12.3.09 750-033205 YR9437            MPC Type 3
   CPU              REV 03   711-035209    YT5857            HMPC PMB 2G
   MIC 0            REV 05   750-033199    YR3295            1X100GE CFP
    PIC 0                    BUILTIN       BUILTIN           1X100GE CFP
     Xcvr 0                  NON-JNPR      X12000187         CFP-100G-SR10
   MIC 1            REV 10   750-033199    YX4518            1X100GE CFP
    PIC 2                    BUILTIN       BUILTIN           1X100GE CFP
     Xcvr 0        REV 01    740-035329    X12J00008         CFP-100G-SR10
  FPC 5            REV 06    750-024884    JW9769            MPC Type 2 3D EQ
  CPU              REV 02    711-028401    JR6158            MPC PMB 2G Proto
  MIC 0            REV 05    750-028387    JR6197            3D 4x 10GE  XFP
```



```
  PIC 0                     BUILTIN       BUILTIN            2x 10GE  XFP
   Xcvr 0         REV 01    740-014289    T07M71112          XFP-10G-SR
   Xcvr 1         REV 02    740-014289    T08L85610          XFP-10G-SR
   PIC 1                    BUILTIN       BUILTIN            2x 10GE  XFP
 MIC 1            REV 22    750-028392    YM0053            3D 20x 1GE(LAN) SFP

  PIC 2                     BUILTIN       BUILTIN            10x 1GE(LAN) SFP
   Xcvr 0         REV 01    740-011613    AM0703S005B        SFP-SX
   Xcvr 1         REV 01    740-011613    E07L01352          SFP-SX
  PIC 3                     BUILTIN       BUILTIN            10x 1GE(LAN) SFP
   Xcvr 5         REV 01    740-013111    6500217            SFP-T
   Xcvr 9         REV 02    740-013111    8499527            SFP-T
 Fan Tray                                                    Left Fan Tray
```

The PIC number for MIC 1 always starts from 2 (even if the first MIC is a 1X100GE CFP or a legacy MIC).

show chassis hardware (QFX3500 Switches)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                                  QFX3500
Routing Engine 0          BUILTIN      BUILTIN           QFX Routing Engine
FPC 0            REV 04   750-044071   BBAR3902          QFX3500-48S4Q-AFI
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           48x 10G-SFP+
  PIC 1                   BUILTIN      BUILTIN           15x 10G-SFP+
  MGMT BRD       REV 02   750-044063   BBAR0398          QFX3500-MGMT-SFP-AFO
    Xcvr 0       REV 01   740-011614   AC0946S0BD1       SFP-LX10
    Xcvr 1       REV 02   740-013111   A281922           SFP-T
Power Supply 0   Rev 04   740-032091   UI00677           JPSU-650W-AC-AFI
Power Supply 1   REV 00   740-041741   VJ00162           JPSU-650W-AC-AFO
Fan Tray 0                                               QFX Fan Tray, Back to
Front Airlfow
Fan Tray 1                                               QFX Fan Tray, Back to
Front Airlfow
Fan Tray 2                                               QFX Fan Tray, Back to
Front Airlfow
```

show chassis hardware detail (QFX3500 Switches)

```
user@switch> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN000TEST5        QFX3500
Routing Engine 0          BUILTIN      BUILTIN           QFX Routing Engine
FPC 0            REV 05   750-036931   EE0823            QFX3500-48S4Q-AFI


  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           48x 10G-SFP+
    Xcvr 0       REV 01   740-030589   S99E270079        SFP+-10G-LPBK
    Xcvr 1       REV 01   740-030589   S9AK450099        SFP+-10G-LPBK
    Xcvr 2       REV 01   740-030589   S99E270078        SFP+-10G-LPBK
    Xcvr 3       REV 01   740-030589   S9AK450098        SFP+-10G-LPBK
    Xcvr 4       REV 01   740-030589   S99E270075        SFP+-10G-LPBK
    Xcvr 5       REV 01   740-030589   S9AK450093        SFP+-10G-LPBK
    Xcvr 6       REV 01   740-030589   S9AK450097        SFP+-10G-LPBK
    Xcvr 7       REV 01   740-030589   S9AK450095        SFP+-10G-LPBK
```



```
    Xcvr 8       REV 01   740-030589   S99E270072       SFP+-10G-LPBK
    Xcvr 9       REV 01   740-030589   S99E270073       SFP+-10G-LPBK
    Xcvr 10      REV 01   740-030589   S99E270080       SFP+-10G-LPBK
    Xcvr 11      REV 01   740-030589   S9AK450169       SFP+-10G-LPBK
    Xcvr 12      REV 01   740-030589   S99E270076       SFP+-10G-LPBK
    Xcvr 13      REV 01   740-030589   S9AK450167       SFP+-10G-LPBK
    Xcvr 14      REV 01   740-030589   S9AK450170       SFP+-10G-LPBK
    Xcvr 15      REV 01   740-030589   S9AK450166       SFP+-10G-LPBK
    Xcvr 16      REV 01   740-030589   S9AK450092       SFP+-10G-LPBK
    Xcvr 17      REV 01   740-030589   S9AK450163       SFP+-10G-LPBK
    Xcvr 18      REV 01   740-030589   S9AK450094       SFP+-10G-LPBK
    Xcvr 19      REV 01   740-030589   S9AK450100       SFP+-10G-LPBK
    Xcvr 20      REV 01   740-030589   S9AK450168       SFP+-10G-LPBK
    Xcvr 21      REV 01   740-030589   S9AK450165       SFP+-10G-LPBK
    Xcvr 22      REV 01   740-030589   S9AK450073       SFP+-10G-LPBK
    Xcvr 23      REV 01   740-030589   S9AK450164       SFP+-10G-LPBK
    Xcvr 24      REV 01   740-030589   S9AK450074       SFP+-10G-LPBK
    Xcvr 25      REV 01   740-030589   SA62270195       SFP+-10G-LPBK
    Xcvr 26      REV 01   740-030589   S9AK450078       SFP+-10G-LPBK
    Xcvr 27      REV 01   740-030589   S9AK450024       SFP+-10G-LPBK
    Xcvr 28    REV 01     740-030589   S9AK450027       SFP+-10G-LPBK
    Xcvr 29      REV 01   740-030589   S9AK450080       SFP+-10G-LPBK
    Xcvr 30      REV 01   740-030589   S9AK450030       SFP+-10G-LPBK
    Xcvr 31      REV 01   740-030589   S9AK450025       SFP+-10G-LPBK
    Xcvr 32      REV 01   740-030589   S9AK450023       SFP+-10G-LPBK
    Xcvr 33      REV 01   740-030589   S9AK450075       SFP+-10G-LPBK
    Xcvr 34      REV 01   740-030589   S9AK450161       SFP+-10G-LPBK
    Xcvr 35      REV 01   740-030589   S9AK450071       SFP+-10G-LPBK
    Xcvr 36      REV 01   740-030589   S9AK450072       SFP+-10G-LPBK
    Xcvr 37      REV 01   740-030589   S9AK450022       SFP+-10G-LPBK
    Xcvr 38      REV 01   740-030589   S9AK450021       SFP+-10G-LPBK
    Xcvr 39      REV 01   740-030589   S9AK450175       SFP+-10G-LPBK
    Xcvr 40      REV 01   740-030589   S9AK450162       SFP+-10G-LPBK
    Xcvr 41      REV 01   740-030589   S99E270074       SFP+-10G-LPBK
    Xcvr 42      REV 01   740-030589   S9AK450174       SFP+-10G-LPBK
    Xcvr 43      REV 01   740-030589   S9AK450077       SFP+-10G-LPBK
    Xcvr 44      REV 01   740-030589   S9AK450076       SFP+-10G-LPBK
    Xcvr 45      REV 01   740-030589   S9AK450026       SFP+-10G-LPBK
    Xcvr 46      REV 01   740-030589   S9AK450079       SFP+-10G-LPBK
    Xcvr 47      REV 01   740-030589   S9AK450029       SFP+-10G-LPBK
  PIC 1                   BUILTIN      BUILTIN          15x 10G-SFP+
    Xcvr 1       REV 01   740-032986   QA170087         QSFP+-40G-SR4
    Xcvr 4       REV 01   740-032986   QA360442         QSFP+-40G-SR4
    Xcvr 8       REV 01   740-032986   QA170091         QSFP+-40G-SR4
    Xcvr 12      REV 01   740-032986   QA170042         QSFP+-40G-SR4
MGMT BRD       REV 08   750-036946   EE0731           QFX3500-MB
Power Supply 0   Rev 04   740-032091   UI00690          QFX PS 650W AC
Power Supply 1   Rev 04   740-032091   UI00679          QFX PS 650W AC
Fan Tray 0                                              QFX Fan Tray
Fan Tray 1                                              QFX Fan Tray
```

show chassis hardware models (QFX3500 Switches)

```
user@switch> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Routing Engine 0          BUILTIN      BUILTIN
FPC 0            REV 02   711-032234   EC4074
Power Supply 0   PSMI 2C  11-d65800    --
```

show chassis hardware clei-models (QFX3500 Switches)

```
user@switch> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Routing Engine 0          BUILTIN
FPC 0            REV 02   711-032234
Power Supply 0   PSMI 2C  11-d65800
```

show chassis hardware clei-models (QFX5100 Switches)

```
user@switch> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Routing Engine 0          BUILTIN      CMMNV10BRA
FPC 0            REV 01   611-053010   CMMNV10BRA
  PIC 0                   BUILTIN      CMMNV10BRA
Power Supply 0   REV 03   740-053352   MUPABHBAA         JPSU-850W-AC-AFO
Power Supply 1   REV 03   740-053352   MUPABHBAA         JPSU-850W-AC-AFO
Fan Tray 0                                               QFX5100-96S-FANAFO
Fan Tray 1                                               QFX5100-96S-FANAFO
Fan Tray 2                                               QFX5100-96S-FANAFO
```

show chassis hardware (QFX10002 Switches)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                12345             QFX10002-36Q
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           RE-QFX10002-36Q
FPC 0            REV 26   750-059497   ACNL1387          QFX10002-36Q
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           36X40G
    Xcvr 0       REV 01   740-038623   MOC15476230389    QSFP+-40G-CU1M
    Xcvr 1       REV 01   740-038623   MOC15476230438    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC15446231917    QSFP+-40G-CU1M
    Xcvr 3       REV 01   740-038623   MOC15446232043    QSFP+-40G-CU1M
    Xcvr 4       REV      740-038624   APF15470032AVB    QSFP+-40G-CU3M
    Xcvr 5       REV      740-038624   APF15470032H15    QSFP+-40G-CU3M
    Xcvr 6       REV      740-038624   APF15470032A9J    QSFP+-40G-CU3M
    Xcvr 7       REV      740-038624   APF15470032AG7    QSFP+-40G-CU3M
    Xcvr 8       REV      740-038624   APF15470032ALD    QSFP+-40G-CU3M
    Xcvr 9       REV 01   740-053203   APF15470071V43    QSFP+-40G-ACU7M
    Xcvr 10      REV 01   740-053203   APF15470071V15    QSFP+-40G-ACU7M
    Xcvr 11      REV 01   740-053203   APF15470071V12    QSFP+-40G-ACU7M
    Xcvr 13      REV      740-038624   APF15470032H1N    QSFP+-40G-CU3M
    Xcvr 18      REV 01   740-053203   APF154800738HW    QSFP+-40G-ACU7M
    Xcvr 19      REV 01   740-038153   MOC12161530041    QSFP+-40G-CU3M
    Xcvr 20      REV 01   740-038153   APF15500034A29    QSFP+-40G-CU3M
    Xcvr 30      REV 01   740-038623   MOC15476230444    QSFP+-40G-CU1M
    Xcvr 31      REV 01   740-032986   QC330038          QSFP+-40G-SR4
    Xcvr 32      REV 01   740-032986   QC290540          QSFP+-40G-SR4
  Mezz           REV 02   711-059316   ACNG9344          QFX10002 36X40G Mezz
Power Supply 0   REV 03   740-054405   1EDN5389293       AC AFO 1600W PSU
Power Supply 1   REV 03   740-054405   1EDN5346300       AC AFO 1600W PSU
Fan Tray 0                                               QFX10002 Fan Tray 0,
Front to Back Airflow - AFO
Fan Tray 1                                               QFX10002 Fan Tray 1,
```

```
Front to Back Airflow - AFO
Fan Tray 2                                                    QFX10002 Fan Tray 2,
Front to Back Airflow - AFO
```

show chassis hardware detail (QFX10002 Switches)

```
user@switch> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                12345             QFX10002-72Q
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN            RE-QFX10002-72Q
ada0    8193 MB  QEMU                QM00001           Virtio Block Disk
ada1    4096 MB  QEMU                QM00002           Virtio Block Disk
ada2     512 MB  QEMU                QM00003           Virtio Block Disk
ada3    1024 MB  QEMU                QM00004           Virtio Block Disk
usb0 (addr 0.1)  UHCI root HUB 0     Intel             uhub0
usb0 (addr 1.1)  EHCI root HUB 0     Intel             uhub1
usb0 (addr 1.2)  product 0x0020 32   vendor 0x8087     uhub2
usb0 (addr 1.3)  Ultra Fit  21891    SanDisk           umass0
FPC 0            REV 05   750-055415 ACAM4724          QFX10002-72Q
CPU                     BUILTIN      BUILTIN           FPC CPU
```

show chassis hardware (QFX10008 and QFX10016 Switches)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DE994             QFX10008
Midplane         REV 28   750-054097   ACPG3671          QFX10008 Midplane
Routing Engine 0          BUILTIN      BUILTIN           Routing Engine
Routing Engine 1          BUILTIN      BUILTIN           Routing Engine
CB 0             REV 03   750-068820   ACPA3224          Control Board
CB 1             REV 03   750-068820   ACPM9059          Control Board
FPC 0            REV 33   750-051354   ACNP4522          ULC-36Q-12Q28
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           36X40G
    Xcvr 0       REV 01   740-038623   MOC16016230802    QSFP+-40G-CU1M
    Xcvr 1       REV 01   740-038623   MOC16016230802    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC16016231080    QSFP+-40G-CU1M
    Xcvr 3       REV 01   740-038623   MOC16016231080    QSFP+-40G-CU1M
    Xcvr 4       REV      740-038624   APF16220038H15    QSFP+-40G-CU3M
    Xcvr 5       REV      740-038624   APF16220038H5M    QSFP+-40G-CU3M
    Xcvr 6       REV      740-038624   APF160600308W8    QSFP+-40G-CU3M
    Xcvr 8       REV      740-038624   APF16210038FFL    QSFP+-40G-CU3M
    Xcvr 9       REV      740-038624   APF16210038F6F    QSFP+-40G-CU3M
    Xcvr 10      REV      740-038624   APF1605003032B    QSFP+-40G-CU3M
    Xcvr 11      REV      740-038624   APF16070030CDB    QSFP+-40G-CU3M
    Xcvr 13      REV      740-038624   APF16210038FEW    QSFP+-40G-CU3M
    Xcvr 15      REV 01   740-052307   APF16100071C1L    QSFP+-40G-ACU7M
    Xcvr 16      REV      740-038625   APF1623005048E    QSFP+-40G-CU5M
    Xcvr 17      REV      740-038625   APF16230050471    QSFP+-40G-CU5M
    Xcvr 18      REV      740-038625   APF1623005044D    QSFP+-40G-CU5M
    Xcvr 19      REV 01   740-052307   APF16100071C30    QSFP+-40G-ACU7M
    Xcvr 20      REV      740-038625   APF16290055004    QSFP+-40G-CU5M
    Xcvr 21      REV 01   740-038153   APF1622003970G    QSFP+-40G-CU3M
    Xcvr 22      REV      740-038624   APF16190036R90    QSFP+-40G-CU3M
    Xcvr 23      REV      740-038624   APF16050030374    QSFP+-40G-CU3M
    Xcvr 24      REV 01   740-038153   APF162400318HC    QSFP+-40G-CU3M
```

```
    Xcvr 30         REV       740-038624   APF1606003097A    QSFP+-40G-CU3M
    Xcvr 31         REV 01    740-052307   APF160500702R9    QSFP+-40G-ACU7M
    Xcvr 32         REV       740-038624   APF16220038GVR    QSFP+-40G-CU3M
FPD Board           REV 07    711-054687   ACPC7158          QFX10000 FPD
Power Supply 0      REV 02    740-049388   1EDL63104D6       QFX10000 AC
Power Supply 1      REV 02    740-049388   1EDL62503XC       QFX10000 AC
Power Supply 2      REV 02    740-049388   1EDL62503XS       QFX10000 AC
Power Supply 3      REV 02    740-049388   1EDL62503T8       QFX10000 AC
Power Supply 4      REV 02    740-049388   1EDL62503TR       QFX10000 AC
Power Supply 5      REV 02    740-049388   1EDL62503T5       QFX10000 AC
FTC 0               REV 15    750-050108   ACPF4227          QFX10000 FTC
FTC 1               REV 15    750-050108   ACPF4228          QFX10000 FTC
Fan Tray 0          REV 09    760-054372   ACNV5506          QFX10008 FHB
Fan Tray 1          REV 09    760-054372   ACNV5365          QFX10008 FHB
SIB 0               REV 27    750-050058   ACPM4212          QFX10008 SIB
SIB 1               REV 27    750-050058   ACPM4253          QFX10008 SIB
SIB 2               REV 27    750-050058   ACPM4174          QFX10008 SIB
SIB 3               REV 27    750-050058   ACPM4191          QFX10008 SIB
SIB 4               REV 27    750-050058   ACPM4216          QFX10008 SIB
SIB 5               REV 27    750-050058   ACPM4286          QFX10008 SIB
```

show chassis hardware detail (QFX10008 and QFX10016 Switches)

```
user@switch> show chassis hardware details
Hardware inventory:
Item                Version   Part number  Serial number     Description
Chassis                                    12345             QFX10008
Midplane            REV 01    750-054097   ACAM1754          QFX10008 Midplane
Routing Engine 0              BUILTIN      BUILTIN           Routing Engine
ada0   8193 MB  QEMU                     QM00001          Virtio Block Disk
ada1   4096 MB  QEMU                     QM00002          Virtio Block Disk
ada2    512 MB  QEMU                     QM00003          Virtio Block Disk
ada3   1024 MB  QEMU                     QM00004          Virtio Block Disk
usb0 (addr 1)  UHCI root HUB 0          Intel            uhub0
usb0 (addr 1)  EHCI root HUB 0          Intel            uhub1
usb0 (addr 2)  product 0x0020 32        vendor 0x8087    uhub2
Routing Engine 1              BUILTIN      BUILTIN           Routing Engine
ada0   8193 MB  QEMU                     QM00001          Virtio Block Disk
ada1   4096 MB  QEMU                     QM00002          Virtio Block Disk
ada2    512 MB  QEMU                     QM00003          Virtio Block Disk
ada3   1024 MB  QEMU                     QM00004          Virtio Block Disk
usb0 (addr 0.1)  UHCI root HUB 0        Intel            uhub0
usb0 (addr 1.1)  EHCI root HUB 0        Intel            uhub1
usb0 (addr 1.2)  product 0x0020 32      vendor 0x8087    uhub2
CB 0                REV 16    750-052688   ACAM7936          Control Board
CB 1                REV 18    750-052688   ACAM7708          Control Board
FPC 0               REV 26    750-051351   ACPJ1372          ULC-60S-6Q Main Board
CPU                         BUILTIN      BUILTIN          FPC CPU
```

show chassis hardware interconnect-device (QFabric Systems)

```
user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item                Version   Part number  Serial number     Description
Chassis             REV 07                                   QFX_olive
Midplane            REV 07    750-021261   BH0208188289      QFX Midplane
CB 0                REV 07    750-021261   BH0208188289      QFXIC08-CB4S
```


show chassis hardware node-device (QFabric Systems)

```
user@switch> show chassis hardware node-device node1
Routing Engine 0    BUILTIN      BUILTIN          QFX Routing Engine
node1            REV 05    711-032234    ED3694            QFX3500-48S4Q-AFI


  CPU                      BUILTIN      BUILTIN          FPC CPU
  PIC 0                    BUILTIN      BUILTIN          48x 10G-SFP+
    Xcvr 8         REV 01    740-030658    AD0946A028B       SFP+-10G-USR
...
```

show chassis hardware (PTX5000 Packet Transport Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN11D1FD7AJA      PTX5000
Midplane         REV 03   711-031896   ABAC5589          Midplane-8S
FPM              REV 08   760-030647   EG1679            Front Panel Display
PDU 0            Rev 05   740-032019   ZE00006           DC Power Dist Unit
  PSM 0          Rev 05   740-032022   ZJ00018           DC 12V Power Supply
  PSM 1          Rev 04   740-032022   ZC00052           DC 12V Power Supply
  PSM 2          Rev 04   740-032022   ZD00051           DC 12V Power Supply
  PSM 3          Rev 05   740-032022   ZJ00060           DC 12V Power Supply
CCG 0            REV 04   750-030653   EG3703            Clock Generator
CCG 1            REV 04   750-030653   EG3698            Clock Generator
Routing Engine 0 REV 05   740-026942   P737A-002231      RE-DUO-2600
Routing Engine 1 REV 06   740-026942   P737A-002438      RE-DUO-2600
CB 0             REV 08   750-030625   EG5519            Control Board
CB 1             REV 08   750-030625   EG5516            Control Board
FPC 0            REV 18   750-036844   EJ3080            FPC
  CPU            REV 12   711-030686   EJ3260            SNG PMB
FPC 2            REV 13   750-036844   EG5065            FPC
  CPU            REV 09   711-030686   EG4082            SNG PMB
  PIC 0          REV 14   750-031913   EG5127            24x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   143363A00240      SFP+-10G-SR
    Xcvr 1       REV 01   740-031981   UK90PZ1           SFP+-10G-LR
    Xcvr 2       REV 01   740-031980   AD1141A04XH       SFP+-10G-SR
    Xcvr 3       REV 01   740-031981   UK90Q46           SFP+-10G-LR
    Xcvr 4       REV 01   740-031980   AD1141A04X4       SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   B11H02560         SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   B11C01589         SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   AD1141A04XF       SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   123363A01094      SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   AK80LKF           SFP+-10G-SR
    Xcvr 12      REV 01   740-031980   183363A01528      SFP+-10G-SR
    Xcvr 14      REV 01   740-031980   193363A01079      SFP+-10G-SR
    Xcvr 15      REV 01   740-031980   AK80MC8           SFP+-10G-SR
    Xcvr 16      REV 01   740-031980   AJC0BHC           SFP+-10G-SR
    Xcvr 19      REV 01   740-021309   J08D26856         SFP+-10G-LR
    Xcvr 21      REV 01   740-031980   AK80KCT           SFP+-10G-SR
    Xcvr 22      REV 01   740-031981   UK90PZL           SFP+-10G-LR
    Xcvr 23      REV 01   740-031980   AK80N1V           SFP+-10G-SR
FPC 3            REV 13   750-036844   EG5074            FPC
  CPU            REV 09   711-030686   EG4064            SNG PMB
  PIC 1          REV 10   750-031903   EG0325            SNG Load
FPC 5            REV 06   750-036844   EH3198            FPC
  CPU
  PIC 0          REV 14   750-031913   EG5134            24x 10GE(LAN) SFP+
```



```
    Xcvr 0        REV 01    740-031980    AK80LBH          SFP+-10G-SR
    Xcvr 1        REV 01    740-031980    B11B03724        SFP+-10G-SR
    Xcvr 2        REV 01    740-031980    AK80FMH          SFP+-10G-SR
    Xcvr 5        REV 01    740-031980    B11J00818        SFP+-10G-SR
    Xcvr 6        REV 01    740-031980    193363A00743     SFP+-10G-SR
    Xcvr 7        REV 01    740-031980    B11B06125        SFP+-10G-SR
    Xcvr 10       REV 01    740-031980    B11H02529        SFP+-10G-SR
    Xcvr 11       REV 01    740-031980    AK80LFB          SFP+-10G-SR
    Xcvr 12       REV 01    740-031980    193363A01061     SFP+-10G-SR
    Xcvr 15       REV 01    740-031980    B11J00687        SFP+-10G-SR
    Xcvr 16       REV 01    740-031980    193363A00738     SFP+-10G-SR
    Xcvr 18       REV 01    740-031980    AK80MQX          SFP+-10G-SR
    Xcvr 19       REV 01    740-021309    J08C17257        SFP+-10G-LR
    Xcvr 22       REV 01    740-031980    B11J00730        SFP+-10G-SR
    Xcvr 23       REV 01    740-031980    AK80KEE          SFP+-10G-SR
  PIC 1           REV 08    750-036710    EG3105           2x 40GE CFP
    Xcvr 0        REV 01    740-034554    B260HLT          CFP-40G-LR4
    Xcvr 1        REV 01    740-034554    B11C02847        CFP-40G-LR4
FPC 6             REV 18    750-036844    EJ4391           FPC
  CPU             REV 12    711-030686    EJ3257           SNG PMB
FPC 7             REV 18    750-036844    EJ4382           FPC
  CPU             REV 12    711-030686    EJ3238           SNG PMB
SPMB 0            REV 10    711-030686    EG5418           SNG PMB
SPMB 1            REV 09    711-030686    EG5373           SNG PMB
SIB 0             REV 07    750-030631    EG4858           SIB-I-8S
SIB 1             REV 07    750-030631    EG4872           SIB-I-8S
SIB 2             REV 07    750-030631    EG4866           SIB-I-8S
SIB 3             REV 07    750-030631    EG6011           SIB-I-8S
SIB 4             REV 07    750-030631    EG4907           SIB-I-8S
SIB 5             REV 07    750-030631    EG4879           SIB-I-8S
SIB 6             REV 07    750-030631    EG4864           SIB-I-8S
SIB 7             REV 07    750-030631    EG4899           SIB-I-8S
SIB 8             REV 07    750-030631    EG4880           SIB-I-8S
Fan Tray 0        REV 04    760-032784    EG1496           Vertical Fan Tray
Fan Tray 1        REV 04    760-030642    EG1335           Horizontal Fan Tray
Fan Tray 2        REV 02    760-030642    ED4952           Horizontal Fan Tray
```

show chassis hardware (PTX5000 Packet Transport Router with AC PSM and PDU)

```
user@host> show chassis hardware
Hardware inventory:
Item              Version   Part number   Serial number    Description
Chassis                                   JN12223A6AJA     PTX5000
Midplane          REV 16    750-035893    ACRA1350         Midplane-8S
FPM               REV 12    760-030647    BBBD5625         Front Panel Display
PDU 0            Rev 01    740-048338    1GB83360005      High Capacity AC WYE PDU

  PSM 0           Rev 01    740-048334    1GB43360074      High Capacity AC PSM
  PSM 1           Rev 01    740-048334    1GB43360001      High Capacity AC PSM
  PSM 2           Rev 01    740-048334    1GB43360104      High Capacity AC PSM
  PSM 3           Rev 01    740-048334    1GB43360042      High Capacity AC PSM
  PSM 4           Rev 01    740-048334    1GB43360068      High Capacity AC PSM
  PSM 5           Rev 01    740-048334    1GB43360080      High Capacity AC PSM
  PSM 6           Rev 01    740-048334    1GB43360046      High Capacity AC PSM
  PSM 7           Rev 01    740-048334    1GB43360100      High Capacity AC PSM
PDU 1            Rev 01    740-048338    1GB83360006      High Capacity AC WYE PDU

  PSM 0           Rev 01    740-048334    1GB43360069      High Capacity AC PSM
  PSM 1           Rev 01    740-048334    1GB43360099      High Capacity AC PSM
  PSM 2           Rev 01    740-048334    1GB43360050      High Capacity AC PSM
  PSM 3           Rev 01    740-048334    1GB43360095      High Capacity AC PSM
```

```
  PSM 4          Rev 01   740-048334   1GB43360101      High Capacity AC PSM
  PSM 5          Rev 01   740-048334   1GB43360075      High Capacity AC PSM
  PSM 6          Rev 01   740-048334   1GB43360047      High Capacity AC PSM
  PSM 7          Rev 01   740-048334   1GB43360019      High Capacity AC PSM
CCG 0            REV 09   750-030653   BBAZ5345         Clock Generator
...
```

show chassis hardware (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                JN1204FC0AJA     PTX5000
Midplane         REV 11   750-035893   ACAB8038         Midplane-8S
FPM              REV 12   760-030647   BBBD5619         Front Panel
Display
PDU 0            Rev 04   740-048336   1GB93470043      High Capacity DC PDU
  PSM 0          Rev 04   740-046988   1GB63500184      High Capacity DC PSM
  PSM 2          Rev 04   740-046988   1GB63500169      High Capacity DC PSM
  PSM 4          Rev 04   740-046988   1GB63500306      High Capacity DC PSM
  PSM 6          Rev 04   740-046988   1GB63500074      High Capacity DC PSM
PDU 1            Rev 04   740-048336   1GB93470045      High Capacity DC PDU
  PSM 1          Rev 04   740-046988   1GB63500193      High Capacity DC PSM
  PSM 3          Rev 04   740-046988   1GB63500143      High Capacity DC PSM
  PSM 5          Rev 04   740-046988   1GB63500146      High Capacity DC PSM
  PSM 7          Rev 04   740-046988   1GB63500192      High Capacity DC PSM
CCG 0            REV 09   750-030653   BBBC1909         Clock Generator
CCG 1            REV 09   750-030653   BBBD2970         Clock Generator
...
```

show chassis hardware clei-models (PTX5000 Packet Transport Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code        FRU model number
FPM              REV 08   760-030647   PROTOXCLEI       CRAFT-PTX5000-S
PDU 0            Rev 05   740-032019   IPUPAHLKAA       PWR-SAN-PDU-DC
  PSM 0          Rev 05   740-032022   IPUPAHNKAA       PSM-PTX-DC-120-S
  PSM 1          Rev 04   740-032022   032022XXXX       PWR-SAN-12-DC
  PSM 2          Rev 04   740-032022   032022XXXX       PWR-SAN-12-DC
  PSM 3          Rev 05   740-032022   IPUPAHNKAA       PSM-PTX-DC-120-S
CCG 0            REV 04   750-030653   PROTOXCLEI       CCG-PTX-S
CCG 1            REV 04   750-030653   PROTOXCLEI       CCG-PTX-S
Routing Engine 0 REV 05   740-026942                    RE-DUO-C2600-16G-S
Routing Engine 1 REV 06   740-026942                    RE-DUO-C2600-16G-S
CB 0             REV 08   750-030625   PROTOXCLEI       CB-PTX-S
CB 1             REV 08   750-030625   PROTOXCLEI       CB-PTX-S
FPC 0            REV 18   750-036844   PROTOXCLEI       FPC-PTX-P1-A
FPC 2            REV 13   750-036844   PROTOXCLEI       FPC-PTX-P1-A
  PIC 0          REV 14   750-031913   PROTOXCLEI       P1-PTX-24-10GE-SFPP
FPC 3            REV 13   750-036844   PROTOXCLEI       FPC-PTX-P1-A
FPC 5
  PIC 0          REV 14   750-031913   PROTOXCLEI       P1-PTX-24-10GE-SFPP
FPC 6            REV 18   750-036844   PROTOXCLEI       FPC-PTX-P1-A
FPC 7            REV 18   750-036844   PROTOXCLEI       FPC-PTX-P1-A
SIB 0            REV 07   750-030631   PROTOXCLEI       SIB-I-PTX5008
SIB 1            REV 07   750-030631   PROTOXCLEI       SIB-I-PTX5008
SIB 2            REV 07   750-030631   PROTOXCLEI       SIB-I-PTX5008
SIB 3            REV 07   750-030631   PROTOXCLEI       SIB-I-PTX5008
SIB 4            REV 07   750-030631   PROTOXCLEI       SIB-I-PTX5008
```



```
SIB 5            REV 07   750-030631   PROTOXCLEI        SIB-I-PTX5008
SIB 6            REV 07   750-030631   PROTOXCLEI        SIB-I-PTX5008
SIB 7            REV 07   750-030631   PROTOXCLEI        SIB-I-PTX5008
SIB 8            REV 07   750-030631   PROTOXCLEI        SIB-I-PTX5008
Fan Tray 1       REV 04   760-030642   PROTOXCLEI        FAN-PTX-H-S
```

show chassis hardware clei-models (PTX5000 Packet Transport Router with AC PSM and PDU)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 16   750-035893   IPMUN00ARA        CHAS-MP-PTX5000-S
FPM              REV 12   760-030647   IPUCA7SCAA        CRAFT-PTX5000-S
PDU 0            Rev 01   740-048338   PROTOACPDU        PDU2-PTX-AC-W
  PSM 0          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 1          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 2          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 3          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 4          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 5          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 6          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 7          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
PDU 1            Rev 01   740-048338   PROTOACPDU        PDU2-PTX-AC-W
  PSM 0          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 1          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 2          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 3          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 4          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 5          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 6          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
  PSM 7          Rev 01   740-048334   PROTOACPSM        PSM2-PTX-AC
CCG 0            REV 09   750-030653   IPUCA7DCAA        CCG-PTX-S
...
```

show chassis hardware clei-models (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 11   750-035893   IPMUN00ARA        CHAS-MP-PTX5000-S
FPM              REV 12   760-030647   IPUCA7SCAA        CRAFT-PTX5000-S
PDU 0            Rev 04   740-048336   IPUPAL7KAA        PDU2-PTX-DC-S
  PSM 0          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
  PSM 2          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
  PSM 4          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
  PSM 6          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
PDU 1            Rev 04   740-048336   IPUPAL7KAA        PDU2-PTX-DC-S
  PSM 1          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
  PSM 3          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
  PSM 5          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
  PSM 7          Rev 04   740-046988   IPUPAL8KAA        PSM2-PTX-DC-S
CCG 0            REV 09   750-030653   IPUCA7DCAA        CCG-PTX-S
CCG 1            REV 09   750-030653   IPUCA7DCAA        CCG-PTX-S
...
```

show chassis hardware detail (PTX5000 Packet Transport Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
```



```
Chassis                                JN11D1FD7AJA      PTX5000
Midplane         REV 03   711-031896   ABAC5589          Midplane-8S
FPM              REV 08   760-030647   EG1679            Front Panel Display
PDU 0            Rev 05   740-032019   ZE00006           DC Power Dist Unit
  PSM 0          Rev 05   740-032022   ZJ00018           DC 12V Power Supply
  PSM 1          Rev 04   740-032022   ZC00052           DC 12V Power Supply
  PSM 2          Rev 04   740-032022   ZD00051           DC 12V Power Supply
  PSM 3          Rev 05   740-032022   ZJ00060           DC 12V Power Supply
CCG 0            REV 04   750-030653   EG3703            Clock Generator
CCG 1            REV 04   750-030653   EG3698            Clock Generator
Routing Engine 0 REV 05   740-026942   P737A-002231      RE-DUO-2600
  ad0    3823 MB  SMART CF             201006190039C02DC02D Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 2011042300CF4C6B4C6B Disk 1
Routing Engine 1 REV 06   740-026942   P737A-002438      RE-DUO-2600
  ad0    3823 MB  SMART CF             20100619053455F055F0 Compact Flash
  ad1   62720 MB  SMART Lite SATA Drive 20110423000AE8E7E8E7 Disk 1
CB 0             REV 08   750-030625   EG5519            Control Board
CB 1             REV 08   750-030625   EG5516            Control Board
FPC 0            REV 18   750-036844   EJ3080            FPC
  CPU            REV 12   711-030686   EJ3260            SNG PMB
FPC 2            REV 13   750-036844   EG5065            FPC
  CPU            REV 09   711-030686   EG4082            SNG PMB
  PIC 0          REV 14   750-031913   EG5127            24x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   143363A00240      SFP+-10G-SR
    Xcvr 1       REV 01   740-031981   UK90PZ1           SFP+-10G-LR
    Xcvr 2       REV 01   740-031980   AD1141A04XH       SFP+-10G-SR
    Xcvr 3       REV 01   740-031981   UK90Q46           SFP+-10G-LR
    Xcvr 4       REV 01   740-031980   AD1141A04X4       SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   B11H02560         SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   B11C01589         SFP+-10G-SR
    Xcvr 8       REV 01   740-031980   AD1141A04XF       SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   123363A01094      SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   AK80LKF           SFP+-10G-SR
    Xcvr 12      REV 01   740-031980   183363A01528      SFP+-10G-SR
    Xcvr 14      REV 01   740-031980   193363A01079      SFP+-10G-SR
    Xcvr 15      REV 01   740-031980   AK80MC8           SFP+-10G-SR
    Xcvr 16      REV 01   740-031980   AJC0BHC           SFP+-10G-SR
    Xcvr 19      REV 01   740-021309   J08D26856         SFP+-10G-LR
    Xcvr 21      REV 01   740-031980   AK80KCT           SFP+-10G-SR
    Xcvr 22      REV 01   740-031981   UK90PZL           SFP+-10G-LR
    Xcvr 23      REV 01   740-031980   AK80N1V           SFP+-10G-SR
FPC 3            REV 13   750-036844   EG5074            FPC
  CPU            REV 09   711-030686   EG4064            SNG PMB
  PIC 1          REV 10   750-031903   EG0325            SNG Load
FPC 5            REV 06   750-036844   EH3198            FPC
  CPU
  PIC 0          REV 14   750-031913   EG5134            24x 10GE(LAN) SFP+
    Xcvr 0       REV 01   740-031980   AK80LBH           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   B11B03724         SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AK80FMH           SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   B11J00818         SFP+-10G-SR
    Xcvr 6       REV 01   740-031980   193363A00743      SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   B11B06125         SFP+-10G-SR
    Xcvr 10      REV 01   740-031980   B11H02529         SFP+-10G-SR
    Xcvr 11      REV 01   740-031980   AK80LFB           SFP+-10G-SR
    Xcvr 12      REV 01   740-031980   193363A01061      SFP+-10G-SR
    Xcvr 15      REV 01   740-031980   B11J00687         SFP+-10G-SR
    Xcvr 16      REV 01   740-031980   193363A00738      SFP+-10G-SR
    Xcvr 18      REV 01   740-031980   AK80MQX           SFP+-10G-SR
    Xcvr 19      REV 01   740-021309   J08C17257         SFP+-10G-LR
    Xcvr 22      REV 01   740-031980   B11J00730         SFP+-10G-SR
```



```
    Xcvr 23         REV 01    740-031980    AK80KEE           SFP+-10G-SR
  PIC 1             REV 08    750-036710    EG3105            2x 40GE CFP
    Xcvr 0          REV 01    740-034554    B260HLT           CFP-40G-LR4
    Xcvr 1          REV 01    740-034554    B11C02847         CFP-40G-LR4
FPC 6               REV 18    750-036844    EJ4391            FPC
  CPU               REV 12    711-030686    EJ3257            SNG PMB
FPC 7               REV 18    750-036844    EJ4382            FPC
  CPU               REV 12    711-030686    EJ3238            SNG PMB
SPMB 0              REV 10    711-030686    EG5418            SNG PMB
SPMB 1              REV 09    711-030686    EG5373            SNG PMB
SIB 0               REV 07    750-030631    EG4858            SIB-I-8S
SIB 1               REV 07    750-030631    EG4872            SIB-I-8S
SIB 2               REV 07    750-030631    EG4866            SIB-I-8S
SIB 3               REV 07    750-030631    EG6011            SIB-I-8S
SIB 4               REV 07    750-030631    EG4907            SIB-I-8S
SIB 5               REV 07    750-030631    EG4879            SIB-I-8S
SIB 6               REV 07    750-030631    EG4864            SIB-I-8S
SIB 7               REV 07    750-030631    EG4899            SIB-I-8S
SIB 8               REV 07    750-030631    EG4880            SIB-I-8S
Fan Tray 0          REV 04    760-032784    EG1496            Vertical Fan Tray
Fan Tray 1          REV 04    760-030642    EG1335            Horizontal Fan Tray
Fan Tray 2          REV 02    760-030642    ED4952            Horizontal Fan Tray
```

show chassis hardware detail (PTX5000 Packet Transport Router with AC PSM and PDU)

```
user@host> show chassis hardware detail
Hardware inventory:
Item                Version   Part number   Serial number     Description
Chassis                                     JN12223A6AJA      PTX5000
Midplane            REV 16    750-035893    ACRA1350          Midplane-8S
FPM                 REV 12    760-030647    BBBD5625          Front Panel Display
PDU 0              Rev 01    740-048338    1GB83360005       High Capacity AC WYE PDU

  PSM 0             Rev 01    740-048334    1GB43360074       High Capacity AC PSM
  PSM 1             Rev 01    740-048334    1GB43360001       High Capacity AC PSM
  PSM 2             Rev 01    740-048334    1GB43360104       High Capacity AC PSM
  PSM 3             Rev 01    740-048334    1GB43360042       High Capacity AC PSM
  PSM 4             Rev 01    740-048334    1GB43360068       High Capacity AC PSM
  PSM 5             Rev 01    740-048334    1GB43360080       High Capacity AC PSM
  PSM 6             Rev 01    740-048334    1GB43360046       High Capacity AC PSM
  PSM 7             Rev 01    740-048334    1GB43360100       High Capacity AC PSM
PDU 1              Rev 01    740-048338    1GB83360006       High Capacity AC WYE PDU

  PSM 0             Rev 01    740-048334    1GB43360069       High Capacity AC PSM
  PSM 1             Rev 01    740-048334    1GB43360099       High Capacity AC PSM
  PSM 2             Rev 01    740-048334    1GB43360050       High Capacity AC PSM
  PSM 3             Rev 01    740-048334    1GB43360095       High Capacity AC PSM
  PSM 4             Rev 01    740-048334    1GB43360101       High Capacity AC PSM
  PSM 5             Rev 01    740-048334    1GB43360075       High Capacity AC PSM
  PSM 6             Rev 01    740-048334    1GB43360047       High Capacity AC PSM
  PSM 7             Rev 01    740-048334    1GB43360019       High Capacity AC PSM
CCG 0               REV 09    750-030653    BBAZ5345          Clock Generator
```

show chassis hardware detail (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis hardware detail
Hardware inventory:
Item                Version   Part number   Serial number     Description
Chassis                                     JN1204FC0AJA      PTX5000
Midplane            REV 11    750-035893    ACAB8038          Midplane-8S
```

```
FPM              REV 12   760-030647   BBBD5619          Front Panel
Display
PDU 0            Rev 04   740-048336   1GB93470043       High Capacity DC PDU
  PSM 0          Rev 04   740-046988   1GB63500184       High Capacity DC PSM
  PSM 2          Rev 04   740-046988   1GB63500169       High Capacity DC PSM
  PSM 4          Rev 04   740-046988   1GB63500306       High Capacity DC PSM
  PSM 6          Rev 04   740-046988   1GB63500074       High Capacity DC PSM
PDU 1            Rev 04   740-048336   1GB93470045       High Capacity DC PDU
  PSM 1          Rev 04   740-046988   1GB63500193       High Capacity DC PSM
  PSM 3          Rev 04   740-046988   1GB63500143       High Capacity DC PSM
  PSM 5          Rev 04   740-046988   1GB63500146       High Capacity DC PSM
  PSM 7          Rev 04   740-046988   1GB63500192       High Capacity DC PSM
CCG 0            REV 09   750-030653   BBBC1909          Clock Generator
CCG 1            REV 09   750-030653   BBBD2970          Clock Generator
...
```

show chassis hardware models (PTX5000 Packet Transport Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
FPM              REV 08   760-030647   EG1679            CRAFT-PTX5000-S
PDU 0            Rev 05   740-032019   ZE00006           PWR-SAN-PDU-DC
  PSM 0          Rev 05   740-032022   ZJ00018           PSM-PTX-DC-120-S
  PSM 1          Rev 04   740-032022   ZC00052           PWR-SAN-12-DC
  PSM 2          Rev 04   740-032022   ZD00051           PWR-SAN-12-DC
  PSM 3          Rev 05   740-032022   ZJ00060           PSM-PTX-DC-120-S
CCG 0            REV 04   750-030653   EG3703            CCG-PTX-S
CCG 1            REV 04   750-030653   EG3698            CCG-PTX-S
Routing Engine 0 REV 05   740-026942   P737A-002231      RE-DUO-C2600-16G-S
Routing Engine 1 REV 06   740-026942   P737A-002438      RE-DUO-C2600-16G-S
CB 0             REV 08   750-030625   EG5519            CB-PTX-S
CB 1             REV 08   750-030625   EG5516            CB-PTX-S
FPC 0            REV 18   750-036844   EJ3080            FPC-PTX-P1-A
FPC 2            REV 13   750-036844   EG5065            FPC-PTX-P1-A
  PIC 0          REV 14   750-031913   EG5127            P1-PTX-24-10GE-SFPP
FPC 3            REV 13   750-036844   EG5074            FPC-PTX-P1-A
FPC 5
  PIC 0          REV 14   750-031913   EG5134            P1-PTX-24-10GE-SFPP
FPC 6            REV 18   750-036844   EJ4391            FPC-PTX-P1-A
FPC 7            REV 18   750-036844   EJ4382            FPC-PTX-P1-A
SIB 0            REV 07   750-030631   EG4858            SIB-I-PTX5008
SIB 1            REV 07   750-030631   EG4872            SIB-I-PTX5008
SIB 2            REV 07   750-030631   EG4866            SIB-I-PTX5008
SIB 3            REV 07   750-030631   EG6011            SIB-I-PTX5008
SIB 4            REV 07   750-030631   EG4907            SIB-I-PTX5008
SIB 5            REV 07   750-030631   EG4879            SIB-I-PTX5008
SIB 6            REV 07   750-030631   EG4864            SIB-I-PTX5008
SIB 7            REV 07   750-030631   EG4899            SIB-I-PTX5008
SIB 8            REV 07   750-030631   EG4880            SIB-I-PTX5008
Fan Tray 1       REV 04   760-030642   EG1335            FAN-PTX-H-S
```

show chassis hardware models (PTX5000 Packet Transport Router with AC PSM and PDU)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 16   750-035893   ACRA1350          CHAS-MP-PTX5000-S
FPM              REV 12   760-030647   BBBD5625          CRAFT-PTX5000-S
PDU 0            Rev 01   740-048338   1GB83360005       PDU2-PTX-AC-W
```



```
  PSM 0          Rev 01   740-048334   1GB43360074       PSM2-PTX-AC
  PSM 1          Rev 01   740-048334   1GB43360001       PSM2-PTX-AC
  PSM 2          Rev 01   740-048334   1GB43360104       PSM2-PTX-AC
  PSM 3          Rev 01   740-048334   1GB43360042       PSM2-PTX-AC
  PSM 4          Rev 01   740-048334   1GB43360068       PSM2-PTX-AC
  PSM 5          Rev 01   740-048334   1GB43360080       PSM2-PTX-AC
  PSM 6          Rev 01   740-048334   1GB43360046       PSM2-PTX-AC
  PSM 7          Rev 01   740-048334   1GB43360100       PSM2-PTX-AC
PDU 1            Rev 01   740-048338   1GB83360006       PDU2-PTX-AC-W
  PSM 0          Rev 01   740-048334   1GB43360069       PSM2-PTX-AC
  PSM 1          Rev 01   740-048334   1GB43360099       PSM2-PTX-AC
  PSM 2          Rev 01   740-048334   1GB43360050       PSM2-PTX-AC
  PSM 3          Rev 01   740-048334   1GB43360095       PSM2-PTX-AC
  PSM 4          Rev 01   740-048334   1GB43360101       PSM2-PTX-AC
  PSM 5          Rev 01   740-048334   1GB43360075       PSM2-PTX-AC
  PSM 6          Rev 01   740-048334   1GB43360047       PSM2-PTX-AC
  PSM 7          Rev 01   740-048334   1GB43360019       PSM2-PTX-AC
CCG 0            REV 09   750-030653   BBAZ5345          CCG-PTX-S
...
```

show chassis hardware models (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 11   750-035893   ACAB8038          CHAS-MP-PTX5000-S
FPM              REV 12   760-030647   BBBD5619          CRAFT-PTX5000-S
PDU 0            Rev 04   740-048336   1GB93470043       PDU2-PTX-DC-S
  PSM 0          Rev 04   740-046988   1GB63500184       PSM2-PTX-DC-S
  PSM 2          Rev 04   740-046988   1GB63500169       PSM2-PTX-DC-S
  PSM 4          Rev 04   740-046988   1GB63500306       PSM2-PTX-DC-S
  PSM 6          Rev 04   740-046988   1GB63500074       PSM2-PTX-DC-S
PDU 1            Rev 04   740-048336   1GB93470045       PDU2-PTX-DC-S
  PSM 1          Rev 04   740-046988   1GB63500193       PSM2-PTX-DC-S
  PSM 3          Rev 04   740-046988   1GB63500143       PSM2-PTX-DC-S
  PSM 5          Rev 04   740-046988   1GB63500146       PSM2-PTX-DC-S
  PSM 7          Rev 04   740-046988   1GB63500192       PSM2-PTX-DC-S
CCG 0            REV 09   750-030653   BBBC1909          CCG-PTX-S
CCG 1            REV 09   750-030653   BBBD2970          CCG-PTX-S
...
```

show chassis hardware extensive (PTX5000 Packet Transport Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
...............
PDU 0            Rev 04   740-032019   UE0003            DC Power Dist Unit
 Jedec Code:   0x7fb0          EEPROM Version:    0x02
 P/N:          740-032019      S/N:               UE0003
 Assembly ID:  0x043d          Assembly Version:  04.00
 Date:         11-29-2010      Assembly Flags:    0x00
 Version:      Rev 04          CLEI Code:         032022XXXX
 ID: DC Power Dist Unit        FRU Model Number:  PWR-SAN-PDU-DC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 3d 04 00 52 65 76 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 31 39 00 00
  Address 0x20: 53 2f 4e 20 55 45 30 30 30 33 00 00 00 1d 0b 07
```

```
  Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
  Address 0x50: 57 52 2d 53 41 4e 2d 50 44 55 2d 44 43 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 a3 ff ff ff ff ff ff ff ff ff ff ff ff
  PSM 0          Rev 04   740-032022   YG00065          DC 12V Power Supply
Module
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          740-032022        S/N:               YG00065
 Assembly ID:  0x0440            Assembly Version:  04.00
 Date:         07-30-2010        Assembly Flags:    0x00
 Version:      Rev 04            CLEI Code:         032022XXXX
 ID: DC 12V Power Supply Module  FRU Model Number:  PWR-SAN-12-DC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 40 04 00 52 65 76 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 32 32 00 00
  Address 0x20: 53 2f 4e 20 59 47 30 30 30 36 35 00 00 1e 07 07
  Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
  Address 0x50: 57 52 2d 53 41 4e 2d 31 32 2d 44 43 20 20 20 20
  Address 0x60: 20 20 20 20 20 20 01 00 ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 0c ff ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware extensive (PTX1000 Packet Transport Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                UNDEFINED         PTX1000
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           RE-PTX1000
FPC 0            REV 06   750-053330   ACAM4850          PTX1000-FPC-P2-BUILTIN
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                  BUILTIN      BUILTIN          288X10GE/72X40GE/24X100GE

    Xcvr 2       REV 01   740-046565   QE240845          QSFP+-40G-SR4
    Xcvr 3       REV 01   740-046565   QE240962          QSFP+-40G-SR4
    Xcvr 5       REV 01   740-032986   ES400LZ           QSFP+-40G-SR4
    Xcvr 12      REV 01   740-054053   QE419452          QSFP+-4X10G-SR
    Xcvr 18      REV 01   740-054053   QE419481          QSFP+-4X10G-SR
    Xcvr 30      REV 01   740-046565   QE440485          QSFP+-40G-SR4
    Xcvr 48      REV 01   740-032986   ES400K3           QSFP+-40G-SR4
    Xcvr 68      REV 01   740-046565   QF2805J3          QSFP+-40G-SR4
  Mezz           REV 05   711-053333   ACAM4282          Mezzanine Board
Power Supply 2   REV 01   740-054405   1EDN4470131       AC AFO 1600W PSU
Power Supply 3   REV 01   740-054405   1EDN4470112       AC AFO 1600W PSU
Fan Tray 0                                             PTX1000 Fan Tray 0, Front
 to Back Airflow - AFO
Fan Tray 1                                             PTX1000 Fan Tray 1, Front
 to Back Airflow - AFO
Fan Tray 2                                             PTX1000 Fan Tray 2, Front
 to Back Airflow - AFO
```

show chassis hardware extensive (PTX5000 with Control Board 2)

```
user@host> show chassis hardware grep CB
CB 0             REV 06   750-055537   ACLZ9541          Control Board 2
CB 1             REV 06   750-055537   ACLY5329          Control Board 2
```


show chassis hardware (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1100FB1AFB      MX480
Midplane         REV 05   710-017414   TR3310            MX480 Midplane
FPM Board        REV 02   710-017254   KG1872            Front Panel Display
PEM 2            Rev 02   740-017343   QCS0812A00N       DC Power Entry Module
PEM 3            Rev 02   740-017343   QCS0812A00U       DC Power Entry Module
Routing Engine 0 REV 07   740-015113   1000740938        RE-S-1300
CB 0             REV 03   710-021523   KF4630            MX SCB
FPC 1            REV 11   750-037207   ZW9726            AS-MCC
  CPU            REV 04   711-038173   ZW4819            AS-MCC PMB
  MIC 0          REV 06   750-037214   ZW3574            AS-MSC
    PIC 0                 BUILTIN      BUILTIN           AS-MSC
  MIC 1          REV 00   750-037211                     AS-MXC
    PIC 2                 BUILTIN      BUILTIN           AS-MXC
```

show chassis hardware extensive (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis hardware extensive
FPC 1            REV 11   750-037207   ZW9726            AS-MCC
 Jedec Code:   0x7fb0          EEPROM Version:    0x02
 P/N:          750-037207      S/N:               ZW9726
 Assembly ID:  0x0b37          Assembly Version:  01.11
 Date:         02-17-2012      Assembly Flags:    0x00
 Version:      REV 11          CLEI Code:         PROTOXCLEI
 ID: AS-MCC          FRU Model Number:  750-037207
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
  Address 0x20: 53 2f 4e 20 5a 57 39 37 32 36 00 00 00 11 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 04   711-038173   ZW4819            AS-MCC-PMB
 Jedec Code:   0x7fb0          EEPROM Version:    0x02
 P/N:          711-038173      S/N:               ZW4819
 Assembly ID:  0x0b38          Assembly Version:  01.04
 Date:         12-30-2011      Assembly Flags:    0x00
 Version:      REV 04
 ID: AS-MCC PMB
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
  Address 0x20: 53 2f 4e 20 5a 57 34 38 31 39 00 00 00 1e 0c 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
  MIC 0          REV 06   750-037214   ZW3574            AS-MSC
```



```
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          750-037214        S/N:               ZW3574
 Assembly ID:  0x0a44            Assembly Version:  01.06
 Date:         02-19-2012        Assembly Flags:    0x00
 Version:      REV 06            CLEI Code:         PROTOXCLEI
 ID: AS-MSC          FRU Model Number:  750-037214
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0a 44 01 06 52 45 56 20 30 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
  Address 0x20: 53 2f 4e 20 5a 57 33 35 37 34 00 00 00 13 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 35 30 2d 30 33 37 32 31 34 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 60 c0 03 e5 f4 00 00 00 00 00 00 00 00
    PIC 0                BUILTIN      BUILTIN            AS-MSC
  MIC 1         REV 00   750-037211                      AS-MXC
 Jedec Code:   0x7fb0            EEPROM Version:    0x01
 P/N:          750-037211
 Assembly ID:  0x0a43            Assembly Version:  01.00
 Date:         255-255-65535     Assembly Flags:    0x00
 Version:      REV 00
 ID: AS-MXC
 Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0a 43 01 00 52 45 56 20 30 30 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 31 00 00
  Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 ff ff ff
  Address 0x30: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff c0 02 e6 6c 7f b0 02 ff 0a 44 01 06
    PIC 2                BUILTIN      BUILTIN            AS-MXC
```

show chassis hardware (ACX5048 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VF3714170810      ACX5048
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           ACX5K Routing Engine
FPC 0            REV 05   650-056267   VF3714170810      ACX5048
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           48x10G-6x40G
    Xcvr 0       REV 02   740-011613   NR2051S           SFP-SX
    Xcvr 33      REV 01   740-030589   SE5N290041        SFP+-10G-LPBK
    Xcvr 35      REV 01   740-030589   SE5N290926        SFP+-10G-LPBK
    Xcvr 37      REV 01   740-030589   SE5N290049        SFP+-10G-LPBK
    Xcvr 39      REV 01   740-030589   SE5N290046        SFP+-10G-LPBK
    Xcvr 48               NON-JNPR     409310098         UNKNOWN
Power Supply 1   REV 03   740-041741   1GA24081097       JPSU-650W-AC-AFO
Fan Tray 0                                              ACX5K Fan Tray 0, Front
 to Back Airflow - AFO
Fan Tray 1                                              ACX5K Fan Tray 1, Front
 to Back Airflow - AFO
```



```
Fan Tray 2                                                ACX5K Fan Tray 2, Front
 to Back Airflow - AFO
Fan Tray 3                                                ACX5K Fan Tray 3, Front
 to Back Airflow - AFO
Fan Tray 4                                                ACX5K Fan Tray 4, Front
 to Back Airflow - AFO
```

show chassis hardware detail (ACX5048 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VF3714170810      ACX5048
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           ACX5K Routing Engine
  ad0      509 MB  QEMU HARDDISK       QM00001           Hard Disk
  ad1     4095 MB  QEMU HARDDISK       QM00002           Hard Disk
  ad2      511 MB  QEMU HARDDISK       QM00003           Hard Disk
  ad3     1023 MB  QEMU HARDDISK       QM00004           Hard Disk
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000     uhub1
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub2
FPC 0            REV 05   650-056267   VF3714170810      ACX5048
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           48x10G-6x40G
    Xcvr 0       REV 02   740-011613   NR2051S           SFP-SX
    Xcvr 33      REV 01   740-030589   SE5N290041        SFP+-10G-LPBK
    Xcvr 35      REV 01   740-030589   SE5N290926        SFP+-10G-LPBK
    Xcvr 37      REV 01   740-030589   SE5N290049        SFP+-10G-LPBK
    Xcvr 39      REV 01   740-030589   SE5N290046        SFP+-10G-LPBK
    Xcvr 48               NON-JNPR     409310098         UNKNOWN
Power Supply 1   REV 03   740-041741   1GA24081097       JPSU-650W-AC-AFO
Fan Tray 0                                              ACX5K Fan Tray 0, Front
 to Back Airflow - AFO
Fan Tray 1                                              ACX5K Fan Tray 1, Front
 to Back Airflow - AFO
Fan Tray 2                                              ACX5K Fan Tray 2, Front
 to Back Airflow - AFO
Fan Tray 3                                              ACX5K Fan Tray 3, Front
 to Back Airflow - AFO
Fan Tray 4                                              ACX5K Fan Tray 4, Front
 to Back Airflow - AFO
```

show chassis hardware clei-models (ACX5048 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Routing Engine 0          BUILTIN      CMMRG00BRA        ACX5048
FPC 0            REV 05   650-056267   CMMRG00BRA        ACX5048
  PIC 0                   BUILTIN      CMMRG00BRA        ACX5048
Power Supply 1   REV 03   740-041741   CMUPABHBAA        JPSU-650W-AC-AFO
Fan Tray 0                                               ACX5K-FAN
Fan Tray 1                                               ACX5K-FAN
Fan Tray 2                                               ACX5K-FAN
Fan Tray 3                                               ACX5K-FAN
Fan Tray 4                                               ACX5K-FAN
```

show chassis hardware models (ACX5048 Router)

```
user@host> show chassis hardware models
```



```
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Routing Engine 0          BUILTIN      BUILTIN           ACX5048
FPC 0            REV 05   650-056267   VF3714170810      ACX5048
  PIC 0                   BUILTIN      BUILTIN           ACX5048
Power Supply 1   REV 03   740-041741   1GA24081097       JPSU-650W-AC-AFO
Fan Tray 0                                               ACX5K-FAN
Fan Tray 1                                               ACX5K-FAN
Fan Tray 2                                               ACX5K-FAN
Fan Tray 3                                               ACX5K-FAN
Fan Tray 4                                               ACX5K-FAN
```

show chassis hardware (ACX5096 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VB3714510139      ACX5096
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           ACX5K Routing Engine
FPC 0            REV 09   650-053391   VB3714510139      ACX5096
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           96x10G-8x40G
    Xcvr 0       REV 01   740-021308   ARS186H           SFP+-10G-SR
    Xcvr 2       REV 01   740-031851   AM1045SUA1G       SFP-SX
    Xcvr 10      REV 02   740-011613   NS11KRP           SFP-SX
    Xcvr 14      REV 01   740-031980   AMC0LKL           SFP+-10G-SR
    Xcvr 20      REV 01   740-021308   ARS18A2           SFP+-10G-SR
    Xcvr 30      REV 02   740-011613   PJ21954           SFP-SX
    Xcvr 35      REV 01   740-031851   PN344LV           SFP-SX
    Xcvr 40      REV 01   740-031851   PLG028R           SFP-SX
    Xcvr 41      REV 01   740-021308   L12D01919         SFP+-10G-SR
    Xcvr 46      REV 01   740-011613   PD91F10           SFP-SX
    Xcvr 64      REV 01   740-031980   AMS0YSS           SFP+-10G-SR
    Xcvr 96      REV 01   740-032986   QE481421          QSFP+-40G-SR4
    Xcvr 99      REV 01   740-032986   QE494942          QSFP+-40G-SR4
    Xcvr 100     REV 01   740-032986   QE494756          QSFP+-40G-SR4
Power Supply 0   REV 01   740-053352   1GD14220106       JPSU-850W-AC-AFO
Power Supply 1   REV 01   740-053352   1GD14220102       JPSU-850W-AC-AFO
Fan Tray 0                                              ACX5K Fan Tray 0, Front
 to Back Airflow - AFO
Fan Tray 1                                              ACX5K Fan Tray 1, Front
 to Back Airflow - AFO
Fan Tray 2                                              ACX5K Fan Tray 2, Front
 to Back Airflow - AFO
```

show chassis hardware detail (ACX5096 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VB3714510139      ACX5096
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           ACX5K Routing Engine
  ad0      509 MB  QEMU HARDDISK       QM00001           Hard Disk
  ad1     4095 MB  QEMU HARDDISK       QM00002           Hard Disk
  ad2      511 MB  QEMU HARDDISK       QM00003           Hard Disk
  ad3     1023 MB  QEMU HARDDISK       QM00004           Hard Disk
  usb0 (addr 1)  product 0x0000 0      vendor 0x0000     uhub1
  usb0 (addr 2)  product 0x0020 32     vendor 0x8087     uhub2
```



```
FPC 0            REV 09   650-053391   VB3714510139      ACX5096
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           96x10G-8x40G
    Xcvr 0       REV 01   740-021308   ARS186H           SFP+-10G-SR
    Xcvr 10      REV 02   740-011613   NS11KRP           SFP-SX
    Xcvr 14      REV 01   740-031980   AMC0LKL           SFP+-10G-SR
    Xcvr 20      REV 01   740-021308   ARS18A2           SFP+-10G-SR
    Xcvr 30      REV 02   740-011613   PJ21954           SFP-SX
    Xcvr 41      REV 01   740-021308   L12D01919         SFP+-10G-SR
    Xcvr 46      REV 01   740-011613   PD91F10           SFP-SX
    Xcvr 64      REV 01   740-031980   AMS0YSS           SFP+-10G-SR
    Xcvr 78      REV 01   740-031851   AM1045SUA1G       SFP-SX
    Xcvr 96      REV 01   740-032986   QE481421          QSFP+-40G-SR4
    Xcvr 99      REV 01   740-032986   QE494942          QSFP+-40G-SR4
    Xcvr 100     REV 01   740-032986   QE494756          QSFP+-40G-SR4
Power Supply 0   REV 01   740-053352   1GD14220106       JPSU-850W-AC-AFO
Power Supply 1   REV 01   740-053352   1GD14220102       JPSU-850W-AC-AFO
Fan Tray 0                                              ACX5K Fan Tray 0, Front
 to Back Airflow - AFO
Fan Tray 1                                              ACX5K Fan Tray 1, Front
 to Back Airflow - AFO
Fan Tray 2                                              ACX5K Fan Tray 2, Front
 to Back Airflow - AFO
```

show chassis hardware clei-models (ACX5096 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Routing Engine 0          BUILTIN      CMMNX10BRA        ACX5096
FPC 0            REV 09   650-053391   CMMNX10BRA        ACX5096
  PIC 0                   BUILTIN      CMMNX10BRA        ACX5096
Power Supply 0   REV 01   740-053352   CMUPACSBAA        JPSU-850W-AC-AFO
Power Supply 1   REV 01   740-053352   CMUPACSBAA        JPSU-850W-AC-AFO
Fan Tray 0                                               ACX5K-FAN
Fan Tray 1                                               ACX5K-FAN
Fan Tray 2                                               ACX5K-FAN
```

show chassis hardware models (ACX5096 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Routing Engine 0          BUILTIN      BUILTIN           ACX5096
FPC 0            REV 09   650-053391   VB3714510139      ACX5096
  PIC 0                   BUILTIN      BUILTIN           ACX5096
Power Supply 0   REV 01   740-053352   1GD14220106       JPSU-850W-AC-AFO
Power Supply 1   REV 01   740-053352   1GD14220102       JPSU-850W-AC-AFO
Fan Tray 0                                               ACX5K-FAN
Fan Tray 1                                               ACX5K-FAN
Fan Tray 2                                               ACX5K-FAN
```

show chassis hardware (ACX500 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                VJ0214510035      ACX500-AC
Midplane         REV 01   650-055932   VJ0214510035      ACX500-AC
Routing Engine            BUILTIN      BUILTIN           Routing Engine
```

```
FEB 0                     BUILTIN      BUILTIN          Forwarding Engine
Processor
FPC 0                     BUILTIN      BUILTIN          FPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN          2x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN          2x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   PMF2Y3C          SFP-SX
      Xcvr 1     REV 01   740-031851   PN342QN          SFP-SX
  MIC 1                   BUILTIN      BUILTIN          4x 1GE(LAN) SFP, RJ45
    PIC 1                 BUILTIN      BUILTIN          4x 1GE(LAN) SFP, RJ45
      Xcvr 0     REV 01   740-011613   PF30K0L          SFP-SX
  MIC 2                   BUILTIN      BUILTIN          MS BUILTIN
    PIC 2                 BUILTIN      BUILTIN          MS BUILTIN
```

show chassis hardware detail (ACX500 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                VJ0214510035     ACX500-AC
Midplane         REV 01   650-055932   VJ0214510035     ACX500-AC
Routing Engine            BUILTIN      BUILTIN          Routing Engine
  da0     3820 MB  USB DISK 2.0                         Nand Flash 0
FEB 0                     BUILTIN      BUILTIN          Forwarding Engine
Processor
FPC 0                     BUILTIN      BUILTIN          FPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN          2x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN          2x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   PMF2Y3C          SFP-SX
      Xcvr 1     REV 01   740-031851   PN342QN          SFP-SX
  MIC 1                   BUILTIN      BUILTIN          4x 1GE(LAN) SFP, RJ45
    PIC 1                 BUILTIN      BUILTIN          4x 1GE(LAN) SFP, RJ45
      Xcvr 0     REV 01   740-011613   PF30K0L          SFP-SX
  MIC 2                   BUILTIN      BUILTIN          MS BUILTIN
    PIC 2                 BUILTIN      BUILTIN          MS BUILTIN
```

show chassis hardware extensive (ACX500 Router)

```
user@host> show chassis hardware extensive
Hardware inventory:
Item             Version  Part number  Serial number    Description
Chassis                                VJ0214510035     ACX500-AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
                                 S/N:               VJ0214510035
 Assembly ID:  0x057c            Assembly Version:  00.00
 Date:         00-00-0000        Assembly Flags:    0x00
 ID: ACX500-AC
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 7f b0 02 ff 05 7c 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: 56 4a 30 32 31 34 35 31 30 30 33 35 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane         REV 01   650-055932   VJ0214510035     ACX500-AC
 Jedec Code:   0x7fb0            EEPROM Version:    0x02
 P/N:          650-055932        S/N:               VJ0214510035
```



```
 Assembly ID:  0x057c           Assembly Version:  01.00
 Date:         12-23-2014       Assembly Flags:    0x00
 Version:      REV 01           CLEI Code:         PROTOXCLEI
 ID: ACX500-AC                  FRU Model Number:  ACX500-AC
 Board Information Record:
  Address 0x00: ad 01 00 80 f0 1c 2d 1b 60 80 ff ff ff ff ff ff
 I2C Hex Data:
  Address 0x00: 7f b0 02 fe 05 7c 01 00 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 36 35 30 2d 30 35 35 39 33 32 00 00
  Address 0x20: 56 4a 30 32 31 34 35 31 30 30 33 35 00 17 0c 07
  Address 0x30: de ff ff ff ad 01 00 80 f0 1c 2d 1b 60 80 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 41
  Address 0x50: 43 58 35 30 30 2d 41 43 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 30 41 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 93 56 4a 30 32 31 34 35 31 30 30 33 35
Routing Engine            BUILTIN      BUILTIN           Routing Engine
  da0    3820 MB  USB DISK 2.0                           Nand Flash 0
FEB 0                     BUILTIN      BUILTIN           Forwarding Engine
Processor
FPC 0                     BUILTIN      BUILTIN           FPC BUILTIN
  MIC 0                   BUILTIN      BUILTIN           2x 1GE(LAN) SFP
 Jedec Code:   0x0000           EEPROM Version:    0x00
 P/N:          BUILTIN          S/N:               BUILTIN
 Assembly ID:  0x0a40           Assembly Version:  00.00
 Date:         00-00-0000       Assembly Flags:    0x00
 ID: 2x 1GE(LAN) SFP
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a 40 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 49 6e 76 61
  Address 0x20: 42 55 49 4c 54 49 4e 00 49 6e 76 61 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 64 20 22 a8 60 af 21 38
    PIC 0                 BUILTIN      BUILTIN           2x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-031851   PMF2Y3C           SFP-SX
      Xcvr 1     REV 01   740-031851   PN342QN           SFP-SX
  MIC 1                   BUILTIN      BUILTIN           4x 1GE(LAN) SFP, RJ45
 Jedec Code:   0x0000           EEPROM Version:    0x00
 P/N:          BUILTIN          S/N:               BUILTIN
 Assembly ID:  0x0aac           Assembly Version:  00.00
 Date:         00-00-0000       Assembly Flags:    0x00
 ID: 4x 1GE(LAN) SFP, RJ45
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a ac 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 49 6e 76 61
  Address 0x20: 42 55 49 4c 54 49 4e 00 49 6e 76 61 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 64 20 b5 c0 60 af 21 38
    PIC 1                 BUILTIN      BUILTIN           4x 1GE(LAN) SFP, RJ45
      Xcvr 0     REV 01   740-011613   PF30K0L           SFP-SX
  MIC 2                   BUILTIN      BUILTIN           MS BUILTIN
 Jedec Code:   0x0000           EEPROM Version:    0x00
```

```
 P/N:           BUILTIN           S/N:               BUILTIN
 Assembly ID:   0x0aaf            Assembly Version:  00.00
 Date:          00-00-0000        Assembly Flags:    0x00
 ID: MS BUILTIN
 Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 I2C Hex Data:
  Address 0x00: 00 00 00 00 0a af 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 49 6e 76 61
  Address 0x20: 42 55 49 4c 54 49 4e 00 49 6e 76 61 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 de ad be ef 64 22 cd 48 60 af 21 38
    PIC 2                BUILTIN      BUILTIN           MS BUILTIN
```

show chassis hardware clei-models (ACX500 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 01   650-055932   PROTOXCLEI        ACX500-AC
Routing Engine            BUILTIN
FEB 0                     BUILTIN
FPC 0                     BUILTIN
```

show chassis hardware models (ACX500 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number     FRU model number
Midplane         REV 01   650-055932   VJ0214510035      ACX500-AC
Routing Engine            BUILTIN      BUILTIN
FEB 0                     BUILTIN      BUILTIN
FPC 0                     BUILTIN      BUILTIN
```


show chassis pic

List of Syntax

Syntax

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Syntax (TX Matrix and TX Matrix Plus Routers)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*
<lcc *number*>

Syntax (MX Series Routers and EX Series Switches)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*
<all-members>
<local>
<member *member-id*>

Syntax (MX104, MX204, MX2010, MX2020, MX10003, and MX2008 3D Universal Edge Routers)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Syntax (PTX Series Packet Transport Router and MX240, MX480, MX960, MX2010, and MX2020 Routers)

show chassis pic transport fpc-slot *slot-number* pic-slot *slot-number*

Syntax (QFX Series)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*
<interconnect-device *name* (fpc-slot *slot-number* | pic-slot *slot-number*)>
<node-device *name* pic-slot *slot-number*>

Syntax (OCX Series)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Syntax (ACX Series Universal Metro Routers)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Syntax (ACX5048 and ACX5096 Routers)

show chassis pic
(fpc-slot *slot-number* | pic-slot *slot-number*)

Syntax (ACX500 Routers)

show chassis pic
(fpc-slot *slot-number* | pic-slot *slot-number*)

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 QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.
Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
transport option introduced in Junos OS Release 16.1R1 for MX Series Routers.
Command introduced in Junos OS Release 17.2 for MX2008 3D Universal Edge Routers.
Command introduced in Junos OS Release 17.2 for PTX10008 Routers.
Command introduced in Junos OS Release 17.3 for MX10003 3D Universal Edge Routers.
Command introduced in Junos OS Release 17.3 for MX150 Router Appliance.
Command introduced in Junos OS Release 17.4 for MX204 3D Universal Edge Routers.
Command introduced in Junos OS Release 18.1R1 for EX9251 switches.
Command introduced in Junos OS Release 18.2 for EX9253 Switches.

Description

Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

Options

fpc-slot *slot-number*—Display information about the PIC in this particular FPC slot:

- On a TX Matrix router, if you specify the number of the T640 router by using the **lcc *number*** option (the recommended method), replace ***slot-number*** with a value from **0** through **7**. Otherwise, replace ***slot-number*** with a value from **0** through **31**.

  Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the **lcc *number*** option (the recommended method), replace ***slot-number*** with a value from **0** through **7**. Otherwise, replace ***slot-number*** with a value from **0** through **31**. For example, the following commands have the same result:

  user@host> **show chassis pic fpc-slot 1 lcc 1 pic-slot 1**
  user@host> **show chassis pic fpc-slot 9 pic-slot 1**

- M120 routers only—Replace ***slot-number*** with a value from **0** through **5**.
- MX80 routers only—Replace ***slot-number*** with a value from **0** through **1**.
- MX104 routers only—Replace ***slot-number*** with a value from **0** through **2**.

- MX240 routers only—Replace ***slot-number*** with a value from **0** through **2**.
- MX480 routers only—Replace ***slot-number*** with a value from **0** through **5**.
- MX960 routers only—Replace ***slot-number*** with a value from **0** through **11**.
- MX2010 routers only—Replace ***slot-number*** with a value from **0** through **9**.
- MX2020 routers only—Replace ***slot-number*** with a value from **0** through **19**.
- MX2008 routers only—Replace ***slot-number*** with a value from **0** through **9**.
- MX10003 routers only—Replace ***slot-number*** with a value from **0** through **1**.
- Other routers—Replace ***slot-number*** with a value from **0** through **7**.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace ***slot-number*** with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace ***slot-number*** with a value from 0 through 9 (switch's member ID).
  - EX8208 switches—Replace ***slot-number*** with a value from 0 through 7 (line card).
  - EX8216 switches—Replace ***slot-number*** with a value from 0 through 15 (line card).
- QFX Series:
  - QFX3500, QFX3600, QFX5100, and OCX Series standalone switches—Replace ***slot-number*** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
  - QFabric systems—Replace ***slot-number*** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

all-members—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

interconnect-device ***name***—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

lcc ***number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

member ***member-id***—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

node-device ***name***—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

pic-slot ***slot-number***—Display information about the PIC in this particular PIC slot. For routers, replace ***slot-number*** with a value from **0** through **3**. For EX3200 and EX4200 switches, replace ***slot-number*** with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace ***slot-number*** with 0. For the QFX3500 standalone switch and the QFabric system, replace ***slot-number*** with 0 or 1.

transport—Display PIC information for optical transport network.

Required Privilege Level

view

Related Documentation

- *request chassis pic*
- show chassis hardware on page 1502
- 100-Gigabit Ethernet Type 4 PIC with CFP Overview on page 441

List of Sample Output

Output Fields

Table 109 on page 1748 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 109: show chassis pic Output Fields

Field Name	Field Description
Type	PIC type. NOTE: On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as **MIC-3D-1STM64-XFP** and with the SONET framing mode, the type is displayed as **MIC-3D-1OC192-XFP**. By default, the 1-port OC192/STM64 MICs displays the type as **MIC-3D-1OC192-XFP**.
Account Layer2 Overhead	(MX Series routers) Indicates whether functionality to count the Layer 2 overhead bytes in the interface statistics at the PIC level is enabled or disabled.
ASIC type	Type of ASIC on the PIC.
State	Status of the PIC. State is displayed only when a PIC is in the slot. • **Online**— PIC is online and running. • **Offline**—PIC is powered down.
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.
Package	(Multiservices PICs only) Services package supported: **Layer-2** or **Layer-3**.
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: **LH**, **LX**, or **SX.**

Table 109: show chassis pic Output Fields (continued)

Field Name	Field Description
PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)	Port-level information for the PIC. • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction • Xcvr Firmware—Transceiver firmware version.

PIC Port Information (MX960 Router Bidirectional Optics)	Port-level information for the PIC.
	• Port—Port number
	• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.
	• Fiber type—Type of fiber. SM is single-mode.
	• Xcvr vendor—Transceiver vendor name.
	• Xcvr vendor part number—Transceiver vendor part number.
	  ◦ BX10-10-km bidirectional optics.
	  ◦ BX40-40-km bidirectional optics.
	  ◦ SFP-LX-40-km SFP optics.
	• Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.

PIC Port Information (Next-Generation SONET/SDH SFP)	Port-level information for the next-generation SONET/SDH SFP PIC. • Port—Port number. • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. • Fiber type—Type of fiber: **SM** (single-mode) or **MM** (multimode). • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.
PIC port information (MX104 router)	Port-level information for the PIC. • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. • Xcvr Firmware—Firmware version of the transceiver.

Table 109: show chassis pic Output Fields (continued)

Field Name	Field Description
Port speed information	Information pertaining to port speed: • Port—Port number. • PFE—Packet Forwarding Engine slot number. • Capable Port Speed—Speed supported by each port.
Multirate Mode	Rate-selectabilty status for the MIC: **Enabled** or **Disabled**.
Channelization	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.
Administrative State	Indicates the administrative state of the PIC. Possible values are: In Service (Default) and Out of Service.
Operational State	Indicates the operational state of the PIC. Possible values are: Normal and Fault.

Sample Output

show chassis pic fpc-slot pic-slot

```
user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
  Type                             10x 1GE(LAN), 1000 BASE
  ASIC type                        H chip
  State                            Online
  PIC version                      1.1
  Uptime                           1 day, 50 minutes, 58 seconds
PIC Port Information:
  Port         Cable           Xcvr              Xcvr Vendor
  Number       Type            Vendor Name       Part Number
  0            GIGE 1000EX     FINISAR CORP.     FTRJ8519P1BNL-J3
  1            GIGE 1000EX     FINISAR CORP.     FTRJ-8519-7D-JUN
```

show chassis pic fpc-slot pic-slot (PIC Offline)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
  State                            Offline
```

show chassis pic fpc-slot pic-slot (FPC Offline)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online
```

show chassis pic fpc-slot pic-slot (FPC Not Present)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty
```


show chassis pic fpc-slot pic-slot (PIC Not Present)

```
user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty
```

show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)

```
user@host>  show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
  Type                             2x G/E IQ, 1000 BASE
  ASIC type                        IQ GE 2 VLAN-TAG FPGA
  State                            Online
  PIC version                 1.16
  Uptime                         3 hours, 3 minutes

PIC Port Information:
  Port        Cable            Xcvr              Xcvr Vendor
  Number      Type             Vendor Name       Part Number
  0           GIGE 1000SX      FINISAR CORP.     FTRJ8519P1BNL-J3
  1           GIGE 1000SX      FINISAR CORP.     FTRJ-8519-7D-JUN
```

show chassis pic fpc-slot pic-slot (MX150)

```
user@host>  show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             Virtual
  State                            Online
  PIC version                  0.0
  Uptime                         7 days, 19 hours, 44 minutes, 40 seconds

PIC port information:
                        Fiber                   Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  10   GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a       0.0

  11   GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a       0.0
```

show chassis pic fpc-slot pic-slot (MX104 Router)

```
user@host>  show chassis pic fpc-slot 1 pic-slot 1
FPC slot 1, PIC slot 1 information:
  Type                             10x 1GE(LAN) -E  SFP
  State                            Online
  PIC version                  1.1
  Uptime                         1 hour, 30 minutes, 59 seconds

PIC port information:
                        Fiber                   Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  3    GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a       0.0

  6    GIGE 1000LX10     SM    FINISAR CORP.      FTLF1318P2BTL-J1  1310 nm   0.0

  8    GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a       0.0
```

```
                     9    GIGE 1000T        n/a   Methode Elec.      SP7041-M1-JN      n/a      0.0
```

show chassis pic fpc-slot pic-slot (MX960 Router with Bidirectional Optics)

```
user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                           10x 1GE(LAN)
  Account Layer2 Overhead        Enabled
  State                          Online
  PIC version                    0.0
  Uptime                         18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:
                          Fiber                        Xcvr vendor
  Port  Cable type        type  Xcvr vendor            part number       Wavelength
  0     SFP-1000BASE-BX10-D SM  SumitomoElectric   SBP6H44-J3-BW-49  1490 nm
  1     SFP-1000BASE-BX10-D SM  SumitomoElectric   SBP6H44-J3-BW-49  1490 nm
  2     SFP-1000BASE-BX10-D SM  SumitomoElectric   SBP6H44-J3-BW-49  1490 nm
  3     SFP-1000BASE-BX10-D SM  OCP                TRXBG1LXDBVM2-JW  1490 nm
  4     SFP-1000BASE-BX10-D SM  OCP                TRXBG1LXDBVM2-JW  1490 nm
  5     SFP-1000BASE-BX10-U SM  SumitomoElectric   SBP6H44-J3-BW-31  1310 nm
  6     SFP-1000BASE-BX10-U SM  SumitomoElectric   SBP6H44-J3-BW-31  1310 nm
  7     SFP-1000BASE-BX10-U SM  OCP                TRXBG1LXDBBMH-J1  1310 nm
  8     SFP-1000BASE-BX10-U SM  OCP                TRXBG1LXDBBMH-J1  1310 nm
  9     SFP-1000BASE-BX10-U SM  SumitomoElectric   SBP6H44-J3-BW-31  1310 nm
```

show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                        1X100GE CFP
  State                       Online
  PIC version                 2.10
  Uptime                      4 minutes, 48 seconds

PIC port information:
        Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor         part number       Wavelength
  0     100GBASE LR4      SM    FINISAR CORP.       FTLC1181RDNS-J3   1310 nm

        Xcvr vendor
        firmware version
        1.8
```

show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```
user@host>show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                             AS-MXC
  State                            Online
  PIC version                  1.0
  Uptime                         11 hours, 18 minutes, 3 seconds
```

show chassis pic fpc-slot pic-slot (MX960 Router with MPC5EQ)

```
user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                             1X100GE CFP2 OTN
```

```
  State                           Online
  PIC version                 0.0
  Uptime                        1 hour, 22 minutes, 42 seconds

PIC port information:
                        Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  0    100GBASE LR4      n/a   Oclaro Inc.        TRB5E20FNF-LF150  1309 nm  1.0
```

show chassis pic fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                            1X100GE DWDM CFP2-ACO
  State                           Online
  PIC version                 1.3
  Uptime                        9 hours, 4 minutes, 43 seconds

PIC port information:
                        Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  0    100G LH           SM    OCLARO             TRB100AJ-01       1528.77 nm -
 1568.36 nm 20.10
```

show chassis pic fpc-slot pic-slot

```
user@host> show chassis pic fpc-slot 1 pic-slot 1

FPC slot 1, PIC slot 1 information:
  Type                            MIC1-MACSEC
  State                           Online
  PIC version                 1.5
  Uptime                        2 hours, 52 minutes, 1 second

PIC port information:
                        Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  8    40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0

  10   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0

Port speed information:

  Port  PFE      Capable Port Speeds
  0      0       4x10GE, 40GE, 100GE
  1      0       4x10GE, 40GE, 100GE
  2      0       4x10GE, 40GE, 100GE
  3      0       4x10GE, 40GE, 100GE
  4      0       4x10GE, 40GE, 100GE
  5      0       4x10GE, 40GE, 100GE
  6      0       4x10GE, 40GE, 100GE
  7      0       4x10GE, 40GE, 100GE
  8      0       4x10GE, 40GE, 100GE
```

```
  9      0       4x10GE, 40GE, 100GE
  10     0       4x10GE, 40GE, 100GE
  11     0       4x10GE, 40GE, 100GE
```

show chassis pic fpc-slot pic-slot (MX10003 Routers)

user@host > **show chassis pic fpc-slot 0 pic-slot 0**

```
FPC slot 0, PIC slot 1 information:
  Type                               MIC1
  State                              Online
  PIC version                    1.5
  Uptime                           13 hours, 54 minutes, 33 seconds

PIC port information:
                         Fiber                   Xcvr vendor       Wave-    Xcvr

  Port Cable type         type  Xcvr vendor        part number       length
Firmware
  0    40GBASE SR4        MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0

  11   40GBASE SR4        MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0

Port speed information:

  Port  PFE      Capable Port Speeds
  0      0       4x10GE, 40GE, 100GE
  1      0       4x10GE, 40GE, 100GE
  2      0       4x10GE, 40GE, 100GE
  3      0       4x10GE, 40GE, 100GE
  4      1       4x10GE, 40GE, 100GE
  5      1       4x10GE, 40GE, 100GE
  6      1       4x10GE, 40GE, 100GE
  7      1       4x10GE, 40GE, 100GE
  8      2       4x10GE, 40GE, 100GE
  9      2       4x10GE, 40GE, 100GE
  10     2       4x10GE, 40GE, 100GE
  11     2       4x10GE, 40GE, 100GE
```

show chassis pic fpc-slot pic-slot (MX204 Routers)

user@host > **show chassis pic fpc-slot 0 pic-slot 0**

```
FPC slot 0, PIC slot 0 information:
  Type                               4XQSFP28 PIC
  State                              Online
  PIC version                    0.0
  Uptime                           2 days, 7 hours, 6 minutes, 10 seconds

PIC port information:
                         Fiber                   Xcvr vendor       Wave-    Xcvr
         JNPR
  Port Cable type         type  Xcvr vendor        part number       length
Firmware     Rev
  0    100GBASE SR4       MM    JUNIPER-FINISAR    FTLC9551REPM-J1   850 nm   0.0
          REV 01
  1    4X10GBASE SR       MM    AVAGO              AFBR-79EEPZ-JU2   850 nm   0.0
          REV 01
  2    100GBASE LR4       SM    JUNIPER-FINISAR    FTLC1151RDPL-J3   1302 nm  0.0
          REV 01
```



```
  3    100GBASE LR4      SM    JUNIPER-FINISAR    FTLC1151RDPL-J3   1302 nm   0.0
          REV 01

Port speed information:

  Port  PFE      Capable Port Speeds
  0      0       4x10GE, 40GE, 100GE
  1      0       4x10GE, 40GE, 100GE
  2      0       4x10GE, 40GE, 100GE
  3      0       4x10GE, 40GE, 100GE
```

show chassis pic fpc-slot pic-slot (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC

```
user@host > show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                              5X100GE DWDM CFP2-ACO
  State                             Online
  PIC version                  1.17
  Uptime                          1 day, 5 hours, 15 minutes, 17 seconds

PIC port information:
                         Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  0    100G LH           SM    MULTILANE SAL      ML4030-ACO-2      1528.77 nm -
 1568.36 nm 1.0
  1    100G LH           SM    MULTILANE SAL      ML4030-ACO-2      1528.77 nm -
 1568.36 nm 1.0
  2    100G LH           SM    JUNIPER-FUJITSU    FIM38500/222      1528.77 nm -
 1568.36 nm 1.16
  3    100G LH           SM    FUJITSU            FIM38500/222      1528.77 nm -
 1568.36 nm 1.16
  4    100G LH           SM    FUJITSU            FIM38500/222      1528.77 nm -
 1568.36 nm 1.16
```

show chassis pic fpc-slot pic-slot (MX480 Router with MPC4E)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                              4x10GE SFPP
  State                             Online
  PIC version                   0.0
  Uptime                          41 seconds

PIC port information:
                         Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor        part number       length
Firmware
  0    10GBASE SR        MM    OPNEXT, INC.       TRS2001EM-0014    850 nm    0.0

  1    10GBASE SR        MM    OPNEXT, INC.       TRS2001EM-0014    850 nm    0.0
```

show chassis pic fpc-slot pic-slot (MX480 router with OTN Interface)

```
user@host> show chassis pci fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                              12X10GE SFPP OTN
```

```
  State                            Online
  PIC version                  0.0
  Uptime                         5 hours, 28 minutes, 23 seconds

PIC port information:
                        Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor         part number        length
Firmware
  0     10GBASE SR       MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0

  1     10GBASE SR       MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0

  2     10GBASE SR       MM    OPNEXT, INC.       TRS2001EM-0014    850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router with OTN Interfaces)

```
user@host> show chassis pic fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type                             2X100GE CFP2 OTN
  State                            Online
  PIC version                  1.9
  Uptime                         3 hours, 56 minutes, 16 seconds

PIC port information:
                        Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor         part number        length
Firmware
  0     100GBASE LR4-D   SM    FUJITSU            FIM37300/222      1310 nm  1.3

  1     100GBASE SR10    MM    AVAGO              AFBR-8420Z        n/a      1.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
  Type                             1X100GE CFP
  Account Layer2 Overhead          Enabled
  State                            Online
  PIC version                      0.0
  Uptime                           14 hours, 51 seconds
```

show chassis pic fpc-slot pic-slot (MX2020 Router)

```
user@host> show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
  Type                             4x 10GE(LAN) SFP+
  Account Layer2 Overhead          Enabled
  State                            Online
  PIC version                  0.0
  Uptime                         1 day, 11 hours, 26 minutes, 36 seconds

PIC port information:
                        Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type        type  Xcvr vendor         part number        length
Firmware
  0     10GBASE SR       MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
```



```
  1     10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0

  2     10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0

  3     10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis pic fpc-slot 18 pic-slot 2
FPC slot 18, PIC slot 2 information:
  Type                             3X40GE QSFPP
  State                            Online
  PIC version                  0.0
  Uptime                         6 minutes, 31 seconds

PIC port information:
                         Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type         type  Xcvr vendor        part number       length
Firmware
  0     40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU2  850 nm   0.0

  1     40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU2  850 nm   0.0

  2     40GBASE SR4       MM    AVAGO              AFBR-79E4Z-D-JU2  850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC6E and OTN MIC)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 0, PIC slot 1 information:
  Type                             24X10GE SFPP OTN
  State                            Online
  PIC version                  1.1
  Uptime                         1 hour, 33 minutes, 59 seconds

PIC port information:
                         Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type         type  Xcvr vendor        part number       length
Firmware
  7     10GBASE SR        MM    SumitomoElectric   SPP5200SR-J6-M    850 nm   0.0

  9     10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0

  12    10GBASE LR        SM    FINISAR CORP.      FTLX1472M3BNL-J3  1310 nm  0.0

  20    10GBASE ZR        SM    FINISAR CORP.      FTLX1871M3BNL-J3  1550 nm  0.0

  21    10GBASE ER        SM    FINISAR CORP.      FTLX1671D3BTL-J4  1550 nm  0.0

  22    10GBASE LR        SM    SOURCEPHOTONICS    SPP10SLREDFCJNP   1310 nm  0.0

  23    10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BNL-J1  1310 nm  0.0
```

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC4E)

```
user@host> show chassis pic fpc-slot 14 pic-slot 0
```

```
FPC slot 14, PIC slot 2 information:
  Type                             4x10GE SFPP
  State                            Online
  PIC version                  0.0
  Uptime                         1 day, 14 hours, 49 minutes, 9 seconds

PIC port information:
                         Fiber                  Xcvr vendor       Wave-    Xcvr

  Port Cable type         type  Xcvr vendor        part number       length
Firmware
  0    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0

  1    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0

  3    10GBASE SR        MM    SumitomoElectric   SPP5100SR-J3      850 nm   0.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
  Type                             1X100GE CFP
  Account Layer2 Overhead          Enabled
  State                            Online
  PIC version                      0.0
  Uptime                           14 hours, 51 seconds
```

show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```
user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
  Type                             100GE SLOT1
  ASIC type                        Brooklyn 100GE FPGA
  State                            Online
  PIC version                  1.3
  Uptime                         10 minutes, 44 seconds

PIC port information:
                          Fiber                   Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number        Wavelength
  0     100GBASE LR4      SM    Opnext Inc.        TRC5E20ENFSF000F   1310 nm
```

show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```
user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:
--------------------------------------------------------------------------
PIC fpc slot 1 pic slot 1 information:
  Type                             4x OC-3 SONET, SMIR
  ASIC type                        D chip
  State                            Online
  PIC version                  1.2
  Uptime                        5 days, 2 hours, 12 minutes, 8 seconds
```

show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```
user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:
--------------------------------------------------------------------------
```



```
FPC slot 8, PIC slot 0 information:
  Type                              1x 10GE(LAN/WAN)
  State                             Online
  Uptime                          2 hours, 46 minutes, 23 seconds

PIC port information:
                          Fiber
  Port  Cable type        type  Xcvr vendor         part number       Wavelength

  0     10GBASE ZR        SM    Opnext Inc.         TRF7061BN-LF150   1550 nm
  0     10GBASE ZR        SM    FINISAR CORP.       FTRX-1811-3-J2    1550 nm
```

show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
   Type                              4x OC-3 1x OC-12 SFP
   ASIC type                         D FPGA
   State                             Online
   PIC version                   1.3
   Uptime                          1 day, 50 minutes, 4 seconds

 PIC port information:
                          Fiber                     Xcvr vendor
   Port  Cable type        type  Xcvr vendor         part number       Wavelength

   0     OC48 short reach  SM    FINISAR CORP.       FTRJ1321P1BTL-J2  1310 nm
   1     OC3 short reach   MM    OCP                 TRPA03MM3BAS-JE   1310 nm
   2     OC3 short reach   MM    OCP                 TRXA03MM3BAS-JW   1310 nm
   3     OC12 inter reach  SM    FINISAR CORP.       FTLF1322P1BTR     1310 nm
```

show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```
user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                              12x T1/E1 CE
  State                             Online
  PIC version                   1.1
  CPU load average                1 percent
  Interrupt load average          0 percent
  Total DRAM size               128 MB
  Memory buffer utilization     100 percent
  Memory heap utilization         4 percent
  Uptime                          1 day, 22 hours, 28 minutes, 12 seconds
  Internal Clock Synchronization   Normal
```

show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                              4x CHOC3 SONET CE SFP
  State                             Online
  PIC version                   1.3
  CPU load average                1 percent
  Interrupt load average          0 percent
  Total DRAM size               128 MB
  Memory buffer utilization      99 percent
  Memory heap utilization         4 percent
  Uptime                          1 day, 22 hours, 55 minutes, 37 seconds
  Internal Clock Synchronization   Normal
```



```
PIC port information:
                            Fiber                     Xcvr vendor
  Port   Cable type         type   Xcvr vendor        part number       Wavelength
  0      OC3 short reach    MM     AVAGO              HFBR-57E0P-JU2    n/a
  1      OC3 short reach    MM     AVAGO              HFBR-57E0P-JU2    n/a
  3      OC3 long reach     SM     OPNEXT INC         TRF5456AVLB314    1310 nm
```

show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                                   MIC-3D-8OC3OC12-4OC48
  State                                  Online
  PIC version                        1.8
  Uptime                               3 days, 22 hours, 3 minutes, 50 seconds

PIC port information:
                            Fiber                     Xcvr vendor
  Port   Cable type         type   Xcvr vendor        part number       Wavelength
  1      OC12 inter reach    SM    FINISAR CORP       FTRJ1322P1BTR-J3  1310 nm
  7      OC12 inter reach    SM    FINISAR CORP       FTRJ1322P1BTR-J3  1310 nm
Multirate Mode                           Enabled
```

show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                                   MIC-3D-8CHOC3-4CHOC12
  State                                  Online
  PIC version                           1.9
  Uptime                                  1 hour, 21 minutes, 24 seconds

PIC port information:
                            Fiber                     Xcvr vendor
  Port   Cable type         type   Xcvr vendor        part number       Wavelength
  0      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  1      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  2      OC12 inter reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J2  1310 nm
  4      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  5      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  6      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  7      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
```

show chassis pic fpc-slot 5 pic-slot 0 (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
FPC slot 5, PIC slot 0 information:
  Type                               MIC-3D-4CHOC3-2CHOC12
  State                              Online
  PIC version                        1.9
  Uptime                             1 hour, 21 minutes

PIC port information:
                            Fiber                     Xcvr vendor
  Port   Cable type         type   Xcvr vendor        part number       Wavelength
  1      OC12 inter reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  2      OC12 inter reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
  3      OC12 short reach   SM     FINISAR CORP.      FTRJ1322P1BTR-J3  1310 nm
```

show chassis pic fpc-slot 1 pic-slot 0 (1-port OC192/STM64 MIC with XFP)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                               MIC-3D-1OC192-XFP
  State                              Online
  PIC version                        1.2
  Uptime                             1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:
                          Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number       Wavelength
  0     OC192 short reach n/a   FINISAR CORP.      FTLX1412M3BCL-J3  1310 nm
```

show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                               MIC-3D-8DS3-E3
  State                              Online
  PIC version                        1.10
  Uptime                             4 days, 1 hour, 29 minutes, 19 seconds
  Channelization Mode                Disabled
```

show chassis pic fpc-slot pic-slot (OTN)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
     Type                               1x10GE(LAN),OTN
     ASIC type                          H chip
     State                              Online
     PIC version                     1.0
     Uptime                          5 minutes, 50 seconds
```

show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
Type 48x 10G-SFP+ Builtin
State Online
Uptime 3 days, 3 hours, 5 minutes, 20 seconds
```

show chassis pic fpc-slot pic-slot (QFX5100 Switches and OCX Series)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                               Unknown Builtin
  State                              Online
  Uptime                           1 day, 17 hours, 5 minutes, 9 seconds
```

show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```
user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type                               16x 40G-GEBuiltin
  State                              Online
  Uptime                           2 hours, 47 minutes, 40 seconds
```


show chassis pic node-device fpc-slot pic-slot (QFabric System)

```
user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
  Type                             48x 10G-SFP+Builtin
  State                            Online
  Uptime                         2 hours, 52 minutes, 37 seconds

PIC port information:
                          Fiber                    Xcvr vendor
  Port  Cable type        type  Xcvr vendor        part number        Wavelength
  0     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  1     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  2     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  3     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  4     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  5     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  6     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  7     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  8     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  9     10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  10    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  11    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  12    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  13    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  14    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  15    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  16    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  17    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  18    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  19    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  20    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  21    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  22    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  23    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  24    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  25    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  26    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  27    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  28    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  29    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  30    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  31    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  32    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  33    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  34    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  35    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  36    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  37    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  38    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  39    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  40    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  41    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  42    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  43    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  44    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  45    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  46    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
  47    10GBASE SR        MM    SumitomoElectric   SPP5101SR-J3       850 nm
```


show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Metro Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                             8x 1GE(LAN) RJ45 Builtin
  State                            Online
  Uptime                         6 days, 2 hours, 51 minutes, 11 seconds
```

show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                             AS-MSC
  State                            Online
  PIC version                    1.6
  Uptime                           11 hours, 17 minutes, 56 seconds
```

show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type                             AS-MXC
  State                            Online
  PIC version                    1.0
  Uptime                           11 hours, 18 minutes, 3 seconds
```

show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)

```
user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic transport fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis pic transport fpc-slot 3 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX5096 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             96x10G-8x40G
  State                            Online
  PIC version                    2.9
  Uptime                           21 hours, 28 minutes, 13 seconds

PIC port information:
                         Fiber                   Xcvr vendor       Wave-    Xcvr

  Port Cable type         type  Xcvr vendor        part number       length
Firmware
  0    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0

  1    10GBASE LR        SM    FINISAR CORP.      FTLX1471D3BCL-J1  1310 nm  0.0

  3    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0
```

```
4     10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
5     10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
6     10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
7     10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
8     10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
9     10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
10    10GBASE SR      MM    OPNEXT, INC.       TRS2001EN-0014     850 nm    0.0
11    10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
12    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
13    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BCL-J1   850 nm    0.0
14    10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
15    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
16    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
17    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
18    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BCL-J1   850 nm    0.0
19    10GBASE LR      SM    FINISAR CORP.      FTLX1471D3BCL-J1   1310 nm   0.0
20    10GBASE LR      SM    FINISAR CORP.      FTLX1471D3BNL-J1   1310 nm   0.0
21    10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
22    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
23    10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
24    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BCL-J1   850 nm    0.0
25    10GBASE USR     MM    FINISAR CORP.      FTLX8570D3BCL-J1   850 nm    0.0
26    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
27    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BCL-J1   850 nm    0.0
28    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
29    GIGE 1000SX     MM    FINISAR CORP.      FTLF8519P3BNL-J1   850 nm    0.0
31    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BNL-J1   850 nm    0.0
32    10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
33    10GBASE SR      MM    FINISAR CORP.      FTLX8571D3BCL-J1   850 nm    0.0
34    10GBASE SR      MM    SumitomoElectric   SPP5100SR-J3       850 nm    0.0
35    10GBASE USR     MM    FINISAR CORP.      FTLX8570D3BCL-J1   850 nm    0.0
```



```
36   10GBASE USR      MM   FINISAR CORP.     FTLX8570D3BCL-J1  850 nm   0.0
37   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BNL-J1  850 nm   0.0
38   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
40   GIGE 1000LX10    SM   FINISAR CORP.     FTLF1318P2BTL-J1  1310 nm  0.0
41   10GBASE LR       SM   OPNEXT,INC        TRS5021EN-S201    1310 nm  0.0
42   10GBASE LR       SM   FINISAR CORP.     FTLX1471D3BCL-J1  1310 nm  0.0
43   10GBASE LR       SM   SumitomoElectric  SPP5100LR-J3      1310 nm  0.0
44   10GBASE LR       SM   SumitomoElectric  SPP5100LR-J3      1310 nm  0.0
45   10GBASE LR       SM   FINISAR CORP.     FTLX1471D3BCL-J1  1310 nm  0.0
46   10GBASE LR       SM   FINISAR CORP.     FTLX1471D3BCL-J1  1310 nm  0.0
47   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BNL-J1  850 nm   0.0
48   10GBASE SR       MM   SumitomoElectric  SPP5100SR-J3      850 nm   0.0
49   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BNL-J1  850 nm   0.0
50   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BNL-J1  850 nm   0.0
51   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
52   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
53   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
54   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
55   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
56   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
57   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
58   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
59   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
60   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
61   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
62   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
63   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
64   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
65   10GBASE SR       MM   FINISAR CORP.     FTLX8571D3BCL-J1  850 nm   0.0
66   10GBASE SR       MM   SumitomoElectric  SPP5200SR-J6-M    850 nm   0.0
```

67	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
68	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
69	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
70	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
71	10GBASE LR	SM	FINISAR CORP.	FTLX1471D3BNL-J1	1310 nm	0.0
72	10GBASE LR	SM	FINISAR CORP.	FTLX1471D3BCL-J1	1310 nm	0.0
73	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
74	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
75	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1	850 nm	0.0
76	10GBASE SR	MM	SumitomoElectric	SPP5100SR-J3	850 nm	0.0
77	10GBASE USR	MM	OPNEXT, INC.	TRS20A0EN-0014	850 nm	0.0
78	10GBASE USR	MM	OPNEXT, INC.	TRS20A0EN-0014	850 nm	0.0
79	10GBASE LRM	MM	OPNEXT INC	TRS5001EN-0014	1310 nm	0.0
80	10GBASE LRM	MM	OPNEXT INC	TRS5001EN-0014	1310 nm	0.0
81	10GBASE USR	MM	OPNEXT, INC.	TRS20A0EN-0014	850 nm	0.0
82	10GBASE USR	MM	OPNEXT, INC.	TRS20A0EN-0014	850 nm	0.0
83	10GBASE USR	MM	OPNEXT, INC.	TRS20A0EN-0014	850 nm	0.0
84	10GBASE USR	MM	OPNEXT, INC.	TRS20A0EN-0014	850 nm	0.0
85	10GBASE LR	SM	OPNEXT,INC	TRS5021EN-S201	1310 nm	0.0
86	10GBASE ER	SM	OPNEXT,INC	TRS7050EN-S201	1550 nm	0.0
87	10GBASE LRM	MM	OPNEXT INC	TRS5001EN-0014	1310 nm	0.0
88	10GBASE LRM	MM	OPNEXT INC	TRS5001EN-0014	1310 nm	0.0
89	10GBASE LRM	MM	OPNEXT INC	TRS5001EN-0014	1310 nm	0.0
90	10GBASE LRM	MM	OPNEXT INC	TRS5001EN-0014	1310 nm	0.0
91	10GBASE USR	MM	FINISAR CORP.	FTLX8570D3BCL-J1	850 nm	0.0
92	10GBASE USR	MM	FINISAR CORP.	FTLX8570D3BCL-J1	850 nm	0.0
93	10GBASE LR	SM	SumitomoElectric	SPP5100LR-J3	1310 nm	0.0
94	10GBASE LR	SM	FINISAR CORP.	FTLX1471D3BNL-J1	1310 nm	0.0
95	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BNL-J1	850 nm	0.0
96	40GBASE SR4	MM	AVAGO	AFBR-79E4Z-D-JU1	850 nm	0.0
97	40GBASE SR4	MM	AVAGO	AFBR-79E4Z-D-JU1	850 nm	0.0

```
98   40GBASE SR4      MM   AVAGO             AFBR-79EQDZ-JU1   850 nm   0.0

99   40GBASE SR4      MM   AVAGO             AFBR-79EQDZ-JU1   850 nm   0.0

100  40GBASE CU 1M    n/a  Molex Inc.        1110409055        n/a      0.0

101  40GBASE CU 1M    n/a  Molex Inc.        1110409055        n/a      0.0

102  40GBASE CU 1M    n/a  Molex Inc.        1110409055        n/a      0.0

103  40GBASE CU 1M    n/a  Molex Inc.        1110409055        n/a      0.0
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX5048 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             96x10G-8x40G
  State                            Online
  PIC version                  2.9
  Uptime                         1 day, 5 hours, 27 minutes, 25 seconds

PIC port information:
                         Fiber                   Xcvr vendor       Wave-    Xcvr

  Port Cable type         type  Xcvr vendor        part number       length
Firmware
  0    10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0

  10   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P3BNL-J1  850 nm   0.0

  14   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0

  20   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BCL-J1  850 nm   0.0

  30   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P2BNL-J1  850 nm   0.0

  41   10GBASE SR        MM    OPNEXT, INC.       TRS2001EN-0014    850 nm   0.0

  46   GIGE 1000SX       MM    FINISAR CORP.      FTLF8519P2BNL-J1  850 nm   0.0

  64   10GBASE SR        MM    FINISAR CORP.      FTLX8571D3BNL-J1  850 nm   0.0

  78   GIGE 1000SX       MM    AVAGO              AFBR-5715PZ-JU2   850 nm   0.0

  96   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0

  99   40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0

  100  40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU1   850 nm   0.0
```

show chassis pic fpc-slot 0 pic-slot 0 (ACX500 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             2x 1GE(LAN) SFP Builtin
  State                            Online
  Uptime                         17 hours, 54 minutes, 45 seconds
```


show chassis pic fpc-slot 0 pic-slot 1 (ACX500 Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                             4x 1GE(LAN) RJ45, SFP Builtin
  State                            Online
  Uptime                         17 hours, 54 minutes, 45 seconds
```

show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)

```
user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic transport fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show chassis pic transport fpc-slot 3 pic-slot 0
Administrative State:      In Service
Operational    State:      Normal
```

show chassis pic fpc-slot 0 pic-slot 0 (EX9251 Switches)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             4XQSFP28 PIC
  State                            Online
  PIC version                  0.0
  Uptime                         1 day, 2 hours, 22 minutes, 3 seconds

PIC port information:
                         Fiber                   Xcvr vendor       Wave-    Xcvr
         JNPR
  Port Cable type        type  Xcvr vendor        part number       length
Firmware     Rev
  0    40GBASE CU 50CM   n/a   Amphenol           601100000         n/a      0.0
          REV 01
  2    40GBASE SR4       MM    AVAGO              AFBR-79EQDZ-JU2   850 nm   0.0
          REV 01

Port speed information:

  Port  PFE      Capable Port Speeds
  0      0       4x10GE, 40GE, 100GE
  1      0       4x10GE, 40GE, 100GE
  2      0       4x10GE, 40GE, 100GE
  3      0       4x10GE, 40GE, 100GE
```

show chassis pic fpc-slot 0 pic-slot 0 (EX9253 Switches)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type                             6xQSFPP
  State                            Online
  PIC version                  0.0
  Uptime                         1 day, 7 minutes, 11 seconds

PIC port information:
                         Fiber                   Xcvr vendor       Wave-    Xcvr
         JNPR
```



```
  Port Cable type        type  Xcvr vendor        part number       length
Firmware     Rev
  0    4X10GBASE SR      MM    AVAGO              AFBR-79EEPZ-JU2   850 nm   0.0
        REV 01

Port speed information:

  Port  PFE      Capable Port Speeds
  0      0       4x10GE, 40GE
  1      0       4x10GE, 40GE
  2      1       4x10GE, 40GE
  3      1       4x10GE, 40GE
  4      2       4x10GE, 40GE
  5      2       4x10GE, 40GE
```

show ethernet-switching redundancy-groups

Syntax

show ethernet-switching redundancy-groups
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs

Release Information

Command introduced in Junos OS Release 13.2.
Command introduced in Junos OS Release 15.1R1 for EX Series switches

Description

Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

Options

redundancy-group-id—(Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.

arp-statistics—(Optional) Count of ARP packets sent and received by the two MC-AE nodes.

nd-statistics—(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.

remote-macs —(Optional) List of remote MAC addresses in the "Installed" state, as learned from the remote MC-AE node.

Required Privilege Level

view

Related Documentation

- *Configuring Multichassis Link Aggregation on EX Series Switches*

List of Sample Output

show ethernet-switching redundancy-groups arp-statistics on page 1772
show ethernet-switching redundancy-groups nd-statistics on page 1772
show ethernet-switching redundancy-groups remote-macs on page 1772
show ethernet-switching redundancy-groups group-id on page 1773

Output Fields

Output fields are listed in the approximate order in which they appear.

Table 110: show ethernet-switching redundancy-groups arp-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.

Table 110: show ethernet-switching redundancy-groups arp-statistics Output Fields (continued)

Field Name	Field Description
MCLAG ARP Statistics Group ID	ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ARP Rx Count From Line	Total number of ARPs received from the Line.
ARP Tx Count To Peer	Total number of ARPs sent to the peer.
ARP Rx Count From Peer	Total number of ARPs received from the peer.
ARP Drop Count received from line	Total number of ARPs sent by the peer that were received.
ARP Drop Count received from peer	Total number of ARPs sent by the peer that were dropped
ARP Install Count	ARP Install Count

Table 111: show ethernet-switching redundancy-groups nd-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ND Statistics Group ID	Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ND Rx Count From Line	Total number of Neighbor Discovery packets received from the Line.
ND Tx Count To Peer	Total number of Neighbor Discovery packets sent to the peer.
ND Rx Count From Peer	Total number of Neighbor Discovery packets received from the peer.
ND Drop Count received from line	Total number of Neighbor Discovery packets sent by the peer that were received.
ND Drop Count received from peer	Total number of Neighbor Discovery packets sent by the peer that were dropped
ND Install Count	ND Install Count

Table 112: show ethernet-switching redundancy-groups remote-macs Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.

Table 112: show ethernet-switching redundancy-groups remote-macs Output Fields (continued)

Field Name	Field Description
Service ID	Service ID (configured at the routing instance level).
Peer-Addr	IP address of the remote peer.
VLAN	Virtual LAN identifier associated with the redundancy group.
MAC	Hardware media access control address associated with the redundancy group.
MCAE-ID	ID number of the MC-AE used by the redundancy group.
Flags	Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected.
Status	Installation state: Installed or Not Installed.

Sample Output

show ethernet-switching redundancy-groups arp-statistics

```
user@host> show ethernet-switching redundancy-groups arp-statistics

Redundancy Group ID : 1       Flags : Local Connect,Remote Connect


 MCLAG ARP Statistics
 Group ID                                : 1
 ARP Rx Count From Line                  : 3493
 ARP Tx Count To Peer                    : 647
 ARP Rx Count From Peer                  : 0
 ARP Install Count                       : 0
 ARP Drop Count received from line       : 2846
 ARP Drop Count received from peer       : 0
```

show ethernet-switching redundancy-groups nd-statistics

```
user@host> show ethernet-switching redundancy-groups nd-statistics

 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                 : 1
ND Rx Count From Line                   : 52
ND Tx Count To Peer                     : 15
ND Rx Count From Peer                   : 39
ND Install Count                        : 34
ND Drop Count received from line        : 37
ND Drop Count received from peer        : 5
```

show ethernet-switching redundancy-groups remote-macs

```
user@host> show ethernet-switching redundancy-groups <redundancy-group-id> remote-macs
```

```
Redundancy Group ID : 1       Flags : Local Connect,Remote Connect


 Service-id  Peer-Addr    VLAN        MAC            MCAE-ID  Subunit  Opcode
Flags   Status
 10         10.3.3.2     100    80:ac:ac:1f:10:a1   1        0        1
0       Installed
```

show ethernet-switching redundancy-groups group-id

`user@host>` **show ethernet-switching redundancy-groups 1**

```
Redundancy Group ID : 1       Flags : Local Connect,Remote Connect
```

show interfaces (Adaptive Services)

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 adaptive services interface.

Options

interface-type—On M Series and T Series routers, the interface type is **sp-** ***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 (Adaptive Services) on page 1779
show interfaces brief (Adaptive Services) on page 1779
show interfaces detail (Adaptive Services) on page 1779
show interfaces extensive (Adaptive Services) on page 1780

Output Fields

Table 113 on page 1774 lists the output fields for the **show interfaces** (adaptive services and redundant adaptive services) command. Output fields are listed in the approximate order in which they appear.

Table 113: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none

Table 113: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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**
Type	Encapsulation being used on the interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source: can be **Internal** or **External**.	All levels
Speed	Speed at which the interface is running.	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 type	Physical interface link type: **Full-Duplex** or **Half-Duplex**.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	**detail extensive** none
Physical info	Information about the physical interface.	**detail 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	MAC address of the hardware.	**detail extensive** none
Alternate link address	Backup address of the link.	**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**

Table 113: Adaptive Services and Redundant Adaptive Services show interfaces 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. NOTE: With static NAT configured as basic NAT44 or destination NAT44 on MX Series routers with MS-MICs and MS-MPCs, the Input bytes field might show 16 more bytes than the Output bytes field. This is caused by the accounting of 16 bytes of the Juniper Forwarding Module cookie. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
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**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant 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. • **Resource errors**—Sum of transmit drops.	**extensive**
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 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. • **MTU errors**—Number of packets larger than the MTU threshold. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number.	**detail extensive** none

Table 113: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical 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.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes generally less than 1 second for the counter to stabilize.	**detail extensive**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, **mpls**.	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**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
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none

Table 113: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Adaptive Services)

```
user@host> show interfaces sp-1/2/0
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : 2006-03-06 11:37:18 PST (00:57:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Input packets : 3057
    Output packets: 3044
    Protocol inet, MTU: 9192
      Flags: Receive-options, Receive-TTL-Exceeded
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.0.0.34, Local: 10.0.0.1
```

show interfaces brief (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 brief
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000

  Logical interface sp-1/2/0.16383
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    inet  10.0.0.1         --> 10.0.0.34
```

show interfaces detail (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 detail
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72, Generation: 30
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2006-03-06 11:37:18 PST (00:57:56 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              125147                    0 bps
   Output bytes  :             1483113                    0 bps
```



```
 Input  packets:                3061                    0 pps
 Output packets:                3048                    0 pps

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
  Traffic statistics:
   Input  bytes  :              125147
   Output bytes  :             1483113
   Input  packets:                3061
   Output packets:                3048
  Local statistics:
   Input  bytes  :              125147
   Output bytes  :             1483113
   Input  packets:                3061
   Output packets:                3048
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Protocol inet, MTU: 9192, Generation: 20, Route table: 1
    Flags: Receive-options, Receive-TTL-Exceeded
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
      Generation: 22
```

show interfaces extensive (Adaptive Services)

```
user@host> show interfaces sp-1/2/0 extensive
Physical interface: sp-1/2/0, Enabled, Physical link is Up
  Interface index: 147, SNMP ifIndex: 72, Generation: 30
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2006-03-06 11:37:18 PST (00:58:40 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              125547                    0 bps
   Output bytes  :             1483353                    0 bps
   Input  packets:                3065                    0 pps
   Output packets:                3052                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Traffic statistics:
     Input  bytes  :              125547
     Output bytes  :             1483353
     Input  packets:                3065
```



```
 Output packets:               3052
Local statistics:
 Input  bytes  :             125547
 Output bytes  :            1483353
 Input  packets:               3065
 Output packets:               3052
Transit statistics:
 Input  bytes  :                  0                    0 bps
 Output bytes  :                  0                    0 bps
 Input  packets:                  0                    0 pps
 Output packets:                  0                    0 pps
Protocol inet, MTU: 9192, Generation: 20, Route table: 1
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
    Generation: 22
```


show interfaces (Aggregated Ethernet)

Syntax

show interfaces ae*number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series, T Series, and MX Series routers only) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.

Options

ae*number*—Display standard information about the specified aggregated Fast Ethernet or 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.

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

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*

List of Sample Output

show interfaces (Aggregated Ethernet) on page 1787
show interfaces brief (Aggregated Ethernet) on page 1788
show interfaces detail (Aggregated Ethernet) on page 1788
show interfaces extensive (Aggregated Ethernet) on page 1789
show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 1790

Output Fields

Table 114 on page 1782 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 114: Aggregated Ethernet show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface and state of the interface.	All levels

Table 114: Aggregated Ethernet show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Enabled	State of the physical 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.	All levels
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
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Minimum links needed	Number of child links that must be operational for the aggregate interface to be operational.	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 “Interfaces Flags” section under *Common Output Fields Description*.	All levels
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up or from up to down. The format is **Last flapped:** ***year-month-day hours:minutes:seconds timezone*** **(*hours:minutes:seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive**
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**

Table 114: Aggregated Ethernet show interfaces 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 and rate, in bps, at which bytes are received on the interface. • **Output bytes**—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • **Input packets**—Number of packets and rate, in pps, at which packets are received on the interface. • **Output packets**—Number of packets and rate, in pps, at which packets are transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of incoming frame aborts and frame check sequence (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 random early detection (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. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**
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), then the cable, the far-end system, or the PIC 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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**

Table 114: Aggregated Ethernet show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Queue counters	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. 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 scenario, 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**
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 of 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 Field" section under *Common Output Fields Description*.	All levels
VLAN-Tag	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
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

Table 114: Aggregated Ethernet show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

Statistics	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.
	• **Bundle**—Information about input and output bundle rates.
	• **Link**—(**detail** and **extensive** only) Information about specific links in the aggregate, including link state and input and output rates.
	• **Adaptive Statistics**—(**extensive** only) Information about adaptive load balancing counter statistics.
	  ◦ **Adaptive Adjusts**—Number of times traffic flow imbalance was corrected by implementation of adaptive load balancing.
	  ◦ **Adaptive Scans**—Number of times the link utilization on each member link of the AE bundle was scanned by for adaptive load balancing
	  ◦ **Adaptive Tolerance**—Tolerance level, in percentage, for load imbalance on link utilization on each member link of the AE bundle.
	  ◦ **Adaptive Updates**—Number of times traffic flow loads have been updated on an AE bundle.
	• **Marker Statistics**—(**detail** and **extensive** only) Information about 802.3ad marker protocol statistics on the specified links.
	  ◦ **Marker Rx**—Number of valid marker protocol data units (PDUs) received on this aggregation port.
	  ◦ **Resp Tx**—Number of marker response PDUs transmitted on this aggregation port.
	  ◦ **Unknown Rx**—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.
	  ◦ **Illegal Rx**—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).
		detail extensive none

LACP info	Link Aggregation Control Protocol (LACP) information for each aggregated interface.
	• **Role** can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **System priority**—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• **System identifier**—Actor or partner system ID, encoded as a MAC address.
	• **Port priority**—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• **Unknown Rx**—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown protocol data unit (PDU), or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.
	• **Port key**—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.	

Table 114: Aggregated Ethernet show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
LACP Statistics	LACP statistics for each aggregated interface. • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received. NOTE: For **LACP Rx** and **LACP Tx**, Packet count is updated only on snmp timer expiry (30 secs).	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	**brief**
Protocol	Protocol family configured on the logical interface. 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
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about protocol family flags. Possible values are described in the “Family Flags Field” section under *Common Output Fields Description*.	**detail extensive** none
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 address flags. 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 interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Aggregated Ethernet)

`user@host>` **show interfaces ae0**

```
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 10.100.1.255
```

show interfaces brief (Aggregated Ethernet)

```
user@host> show interfaces ae0 brief
Physical interface: ae0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384

  Logical interface ae0.0
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    inet  203.0.113.2/24
```

show interfaces detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
  Last flapped   : Never
  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
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                 7375                 7375                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0
```



```
    3 network-cont              2268                2268                    0

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Link:
      fe-0/1/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/3.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      fe-0/1/0.0                 0           0            0            0
      fe-0/1/2.0                 0           0            0            0
      fe-0/1/3.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 37, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
        Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,

        Generation: 49
```

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   60                   0 bps
   Output bytes  :                    0                   0 bps
   Input  packets:                    1                   0 pps
   Output packets:                    0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort               7375                7375                    0

    1 expedited-fo                 0                   0                    0

    2 assured-forw                 0                   0                    0
```



```
    3 network-cont              2268                2268                    0

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             1          0            60            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      fe-0/1/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      fe-0/1/3.0
        Input :             1          0            60            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port     Port  Port

                            priority         identifier  priority  number   key

      fe-1/0/3.0     Actor        127  00:00:5e:00:53:85       127       2     1

      fe-1/0/3.0   Partner        127  00:00:5e:00:53:c3       127       1     1

    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      fe-1/0/3.0              3188        3186            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      fe-0/1/0.0                 0           0            0            0
      fe-0/1/2.0                 0           0            0            0
      fe-0/1/3.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 37, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,

        Generation: 49
```

show interfaces extensive (Aggregated Ethernet with VLAN Stacking)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 48, Generation: 186
  Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5e:00:53:3f, Hardware address: 00:00:5e:00:53:3f
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             2406875                40152 bps
   Output bytes  :             1124470                22056 bps
   Input  packets:                5307                    5 pps
```



```
  Output packets:               13295                  21 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, Giants: 0, Policed discards:
0, Resource errors: 0
 Output errors:
   Carrier transitions: 0, Errors: 0, Drops: 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                   0               859777                    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              1897615                    0

   1 expedited-fo                  0                    0                    0

   2 assured-forw                  0                    0                    0

   3 network-cont                  0               662505                    0


 Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
   Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ]  Encapsulation: VLAN-VPLS
   Statistics        Packets        pps         Bytes          bps
   Bundle:
       Input :           289          0         25685          376
       Output:          1698          4        130375         3096
   Link:
     ge-1/2/0.451
       Input :           289          0         25685          376
       Output:             0          0             0            0
     ge-1/2/1.451
       Input :             0          0             0            0
       Output:          1698          4        130375         3096
   Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
     ge-1/2/0.451               0           0            0            0
     ge-1/2/1.451               0           0            0            0
   Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
     Flags: Is-Primary

 Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)
   Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ]  Encapsulation: VLAN-VPLS
   Statistics        Packets        pps         Bytes          bps
   Bundle:
       Input :           293          1         26003         1072
       Output:          1694          3        130057         2400
   Link:
```



```
      ge-1/2/0.452
        Input :          293          1        26003        1072
        Output:         1694          3       130057        2400
      ge-1/2/1.452
        Input :            0          0            0           0
        Output:            0          0            0           0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/2/0.452               0           0            0            0
      ge-1/2/1.452               0           0            0            0
    Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
      Flags: None
...
```

show interfaces demux0 (Demux Interfaces)

Syntax

show interfaces demux0.*logical-interface-number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 9.0.

Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

Options

none—Display standard information about the specified demux 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

Related Documentation

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show interfaces demux0 (Demux) on page 1799
show interfaces demux0 (PPPoE over Aggregated Ethernet) on page 1800
show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links) on page 1801
show interfaces demux0 (ACI Interface Set Configured) on page 1801

Output Fields

Table 115 on page 1793 lists the output fields for the **show interfaces** demux0 (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	**brief detail extensive** none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**brief detail extensive** none
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	**brief detail extensive** none
Physical link	Status of the physical link (**Up** or **Down**).	**detail extensive** none
Admin	Administrative state of the interface (**Up** or **Down**).	**terse**
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
Link	Status of the physical link (**Up** or **Down**).	**terse**
Targeting summary	Status of aggregated Ethernet links that are configured with targeted distribution (**primary** or **backup**)	**extensive**
Bandwidth	Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.	**extensive**
Proto	Protocol family configured on the interface.	**terse**
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**
Type	Type of interface. **Software-Pseudo** indicates a standard software interface with no associated hardware device.	**brief detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**brief detail extensive**
MTU	Maximum transmission unit size on the physical interface.	**brief detail extensive**
Clocking	Reference clock source: **Internal** (1) or **External** (2).	**brief detail extensive**
Speed	Speed at which the interface is running.	**brief detail extensive**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	**brief detail extensive** none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	**brief detail extensive** none
Link type	Data transmission type.	**detail extensive** none

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	**detail extensive** none
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Alternate link address	Backup address of the link.	**detail extensive**
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
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. • **IPv6 transit statistics**—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined. ◦ **Input bytes**—Number of bytes received on the interface ◦ **Output bytes**—Number of bytes transmitted on the interface. ◦ **Input packets**—Number of packets received on the interface. ◦ **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface whose definitions are as follows: • **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. • **Giants**—Number of frames received that are larger than the giant packet 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. • **Resource errors**—Sum of transmit drops.	**extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none
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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Output Rate	Output rate in bps and pps.	none
Logical Interface		
Logical interface	Name of the logical interface.	**brief detail extensive** none
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**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	**brief detail extensive** none
Encapsulation	Encapsulation on the logical interface.	**brief extensive** none

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying demux interface to create dynamic VLAN subscriber interfaces based on ACI information.	**brief detail extensive** none
Demux	Specific IP demultiplexing (demux) values: • **Underlying interface**—The underlying interface that the demux interface uses. • **Index**—Index number of the logical interface. • **Family**—Protocol family configured on the logical interface. • **Source prefixes, total**—Total number of source prefixes for the underlying interface. • **Destination prefixes, total**—Total number of destination prefixes for the underlying interface. • **Prefix**—**inet** family prefix.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface.	**brief**

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.
	• **IPv6 transit statistics**—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.
	NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.
	  ◦ **Input bytes**—Number of bytes received on the interface.
	  ◦ **Output bytes**—Number of bytes transmitted on the interface.
	  ◦ **Input packets**—Number of packets received on the interface.
	  ◦ **Output packets**—Number of packets transmitted on the interface.
		detail extensive

Local statistics	Number of transit bytes and packets received and transmitted on the local 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.	**detail extensive**

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 Transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input packets	Number of packets received on the interface.	none
Output packets	Number of packets transmitted on the interface.	none
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
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**
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none
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
Destination	IP address of the remote side of the connection.	**detail extensive statistics** none

Table 115: show interfaces demux0 (Demux Interfaces) Output Fields (continued)

Field Name	Field Description	Level of Output
Local	IP address of the logical interface.	**detail extensive terse** none
Remote	IP address of the remote interface.	**terse**
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	**detail extensive** none
Dynamic-profile	Name of the PPPoE dynamic profile assigned to the underlying interface.	**detail extensive** none
Service Name Table	Name of the PPPoE service name table assigned to the PPPoE underlying interface.	**detail extensive** none
Max Sessions	Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.	**detail extensive** none
Duplicate Protection	State of duplicate protection: **On** or **Off**. Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface.	**detail extensive** none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**detail extensive** none

Sample Output

show interfaces demux0 (Demux)

```
user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 79, Generation: 129
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  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, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Demux:
     Underlying interface: ge-2/0/1.0 (Index 74)
   Family Inet Source prefixes, total 1
   Prefix: 203.0.113/24
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                1554
     Input  packets:                   0
     Output packets:                  37
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                1554
     Input  packets:                   0
     Output packets:                  37
    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: 395, 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: 203.0.113.13, Broadcast: 203.0.113.255,

         Generation: 434
```

show interfaces demux0 (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]
```

```
    Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 199)
    Link:
      ge-1/0/0
      ge-1/1/0
    Input packets : 0
    Output packets: 0
    Protocol pppoe
      Dynamic Profile: pppoe-profile,
      Service Name Table: service-table1,
      Max Sessions: 100, Duplicate Protection: On,
      Direct Connect: Off,
      AC Name: pppoe-server-1
```

show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links)

user@host> **show interfaces demux0.1073741824 extensive**

```
  Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Demux:
      Underlying interface: ae0 (Index 201)
    Link:
      ge-1/0/0
      ge-1/1/0
      ge-2/0/7
      ge-2/0/8
    Targeting summary:
      ge-1/1/0, primary, Physical link is Up
      ge-2/0/8, backup, Physical link is Up
    Bandwidth: 1000mbps
```

show interfaces demux0 (ACI Interface Set Configured)

user@host> **show interfaces demux0.1073741827**

```
  Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1802 0x8100.302 ]  Encapsulation:
 ENET2
    Demux: Source Family Inet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    Demux:
      Underlying interface: ge-1/0/0 (Index 138)
    Input packets : 18
    Output packets: 16
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re, Unnumbered
      Donor interface: lo0.0 (Index 322)
      Preferred source address: 203.0.113.202
      Addresses, Flags: Primary Is-Default Is-Primary
        Local: 203.0.113.119
    Protocol pppoe
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
```


show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)

Syntax show interfaces diagnostics optics *interface-name*

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for PTX Series routers.

Description Display diagnostics data, warnings, and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, or Virtual Chassis port interfaces.

Options *interface-name*—Interface name. For example:

- ge-*fpc/pic/port*
- et-*fpc/pic/port*
- et-*fpc/pic/port*:*channel*
- xe-*fpc/pic/port*
- vcp-*fpc/pic/port*

Additional Information The transceivers are polled in 1-second intervals for diagnostics data, warnings, and alarms. The alarms do not cause the links to go down or the LEDs to change color, nor generate SNMP traps. Changes in alarm and warning status generate system log messages.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transceiver vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a device is not working.

NOTE: Some transceivers do not support all optical diagnostics features described in the output fields.

If optics measures transmit or receive power as zero, then, the measured power is displayed as **0.000 mW / - Inf dBm**

You can configure the P2-10G-40G-QSFPP PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode. When the PIC is in 40-Gigabit Ethernet mode, you must execute the **show interfaces diagnostics optics et-*fpc/pic/port*** command. The output of this command displays the diagnostic optics information about the corresponding 40-Gigabit Ethernet port of the PIC. However, when the PIC is in 10-Gigabit Ethernet mode, you must execute the **show interfaces diagnostics optics et-*fpc/pic/port*:*channel*** command. The output of this command displays the diagnostic optics information about the corresponding 10-Gigabit Ethernet port of the PIC. For information about the P2-10G-40G-QSFPP PIC, see "P2-10G-40G-QSFPP PIC Overview" on page 409.

Required Privilege Level view

Related Documentation

- *Determining Transceiver Support and Specifications*

List of Sample Output

Output Fields Table 116 on page 1804 lists the output fields for the **show interfaces diagnostics optics** command for DWDM and DWDM OTN PICs. Output fields are listed in the approximate order in which they appear.

Table 116: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the **LsPOWMON** pin in hardware.
Receiver signal average optical power	Average received optical power, in mW and dBm. This indicator is a software equivalent to the **RxPOWMON** pin in hardware. Average optical power is vendor-specific.

Table 116: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs (continued)

Field Name	Field Description
Laser end-of-life alarm	Laser end-of-life alarm: **On** or **Off**.
Laser wavelength alarm	Laser wavelength alarm: **On** or **Off**.
Laser bias current alarm	Laser bias current alarm: **On** or **Off**.
Laser temperature alarm	Laser temperature alarm: **On** or **Off**.
Laser power alarm	Laser power alarm: **On** or **Off**.
Modulator temperature alarm	Modulator temperature alarm: **On** or **Off**. Transceivers from some vendors do not support this field.
Modulator bias alarm	Modulator bias alarm: **On** or **Off**.
Tx multiplexer FIFO error alarm	Transmit multiplexer first in, first out (FIFO) error alarm: **On** or **Off**.
Tx loss of PLL lock alarm	Transmit loss of phase-locked loop (PLL) lock alarm: **On** or **Off**.
Rx loss of average optical power alarm	Receive loss of average optical power alarm: **On** or **Off**.
Rx loss of AC power alarm	Receive loss of AC power alarm: **On** or **Off**. Transceivers from some vendors do not support this field.
Rx loss of PLL lock alarm	Receive loss of phase-locked loop (PLL) lock alarm: **On** or **Off**.

Table 117 on page 1805 lists the output fields for the **show interfaces diagnostics optics** command when the router is operating with bidirectional SFP optics. Output fields are listed in the approximate order in which they appear.

Table 117: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).

Table 117: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (continued)

Field Name	Field Description
Module temperature	Temperature of the optics module, in Celsius and Fahrenheit.
Module voltage	Internally measured module voltage.
Receiver signal average optical power	Average received optical power, in mW and dBm.
Wavelength Channel number	Wavelength channel number set in the optics module.
Wavelength setpoint	Wavelength set in the optics module.
Tx Dither	Transmit dither status. Displays whether transmit dither is enabled or disabled.
Frequency Error	Frequency error reported from optics module.
Wavelength Error	Wavelength error reported from optics module.
Laser bias current high alarm	Laser bias power setting high alarm. Displays **on** or **off**.
Laser bias current low alarm	Laser bias power setting low alarm. Displays **on** or **off**.
Laser bias current high warning	Laser bias power setting high warning. Displays **on** or **off**.
Laser bias current low warning	Laser bias power setting low warning. Displays **on** or **off**.
Laser output power high alarm	Laser output power high alarm. Displays **on** or **off**.
Laser output power low alarm	Laser output power low alarm. Displays **on** or **off**.
Laser output power high warning	Laser output power high warning. Displays **on** or **off**.
Laser output power low warning	Laser output power low warning. Displays **on** or **off**.
Module temperature high alarm	Module temperature high alarm. Displays **on** or **off**.
Module temperature low alarm	Module temperature low alarm. Displays **on** or **off**.

Table 117: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (continued)

Field Name	Field Description
Module temperature high warning	Module temperature high warning. Displays **on** or **off**.
Module temperature low warning	Module temperature low warning. Displays **on** or **off**.
Module voltage high alarm	Module voltage high alarm. Displays **on** or **off**.
Module voltage low alarm	Module voltage low alarm. Displays **on** or **off**.
Module voltage high warning	Module voltage high warning. Displays **on** or **off**.
Module voltage low warning	Module voltage high warning. Displays **on** or **off**.
Laser rx power high alarm	Receive laser power high alarm. Displays **on** or **off**.
Laser rx power low alarm	Receive laser power low alarm. Displays **on** or **off**.
Laser rx power high warning	Receive laser power high warning. Displays **on** or **off**.
Laser rx power low warning	Receive laser power low warning. Displays **on** or **off**.
TEC fault alarm	TEC fault alarm. Displays **on** or **off**.
Wavelength unlocked alarm	Wavelength unlocked alarm. Displays **on** or **off**.
TxTune	Optical transmit side status. Displays whether optical transmit side is not ready due to tuning.
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm: **70.000 mA.**
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm: **0.0002 mA.**
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning: **65.000 mA.**
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning: **0.0002 mA.**

Table 117: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (continued)

Field Name	Field Description
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm: **1.0000 mW** or **0.00 dBm.**
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm: **0.0560 mW** or **-12.52 dBm.**
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning: **0.6300 mW** or **-2.01 dBm.**
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning: **0.0890 mW** or **-10.51 dBm.**
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm: **100**° C or **212**° F.
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm: **-50**° C or **-58**° F.
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning: **95** ° C or **203** ° F.
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning: **-48**° C or **-54**° F.
Module voltage high alarm threshold	Module voltage high alarm threshold: **3.700 v**.
Module voltage low alarm threshold	Module voltage low alarm threshold: **2.900 v**.
Module voltage high warning threshold	Module voltage high warning threshold: **3.7600 v**.
Module voltage low warning threshold	Module voltage low warning threshold: **3.000 v**.
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm: **1.9953 mW** or **3.00 dBm.**
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm: **0.0001 mW** or **-40.00 dBm.**
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning: **1.0000 mW** or **0.00 dBm.**
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning: **0.0010 mW** or **-30.00 dBm.**

Table 118 on page 1809 lists the output fields for the **show interfaces diagnostics optics** command for SFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 118: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Measured laser bias current in uA.
Laser output power	Measured laser output power in mW.
Module temperature	Internally measured module temperature.
Module voltage	Internally measured module voltage.
Laser rx power	Measured receive optical power in mW.
Laser bias current high alarm	Laser bias current high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser bias current low alarm	Laser bias current low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power high alarm	Laser output power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power low alarm	Laser output power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp high alarm	Module temperature high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp low alarm	Module temperature low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power high alarm	Laser receive power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power low alarm	Laser receive power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser bias current high warning	Laser bias current high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current low warning	Laser bias current low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power high warning	Laser output power high warning: **On** or **Off**. Warning ranges are vendor-specific.

Table 118: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (continued)

Field Name	Field Description
Laser output power low warning	Laser output power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature high warning	Module temperature high warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature low warning	Module temperature low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power high warning	Laser receive power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power low warning	Laser receive power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current high alarm threshold	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current low alarm threshold	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current high warning threshold	Laser bias current high warning threshold. Warning ranges are vendor-specific.
Laser bias current low warning threshold	Laser bias current low warning threshold. Warning ranges are vendor-specific.
Laser output power high alarm threshold	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power low alarm threshold	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power high warning threshold	Laser output power high warning threshold. Warning ranges are vendor-specific.
Laser output power low warning threshold	Laser output power low warning threshold. Warning ranges are vendor-specific.
Module temperature high alarm threshold	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature low alarm threshold	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature high warning threshold	Module temperature high warning threshold. Warning ranges are vendor-specific.

Table 118: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (continued)

Field Name	Field Description
Module temperature low warning threshold	Module temperature low warning threshold. Warning ranges are vendor-specific.
Module voltage high alarm threshold	Module voltage high alarm threshold. Alarm ranges are vendor-specific.
Module voltage low alarm threshold	Module voltage low alarm threshold. Alarm ranges are vendor-specific.
Module voltage high warning threshold	Module voltage high warning threshold. Warning ranges are vendor-specific.
Module voltage low warning threshold	Module voltage low warning threshold. Warning ranges are vendor-specific.
Laser rx power high alarm threshold	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power low alarm threshold	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power high warning threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
Laser rx power high low threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.

Table 119 on page 1811 lists the output fields for the **show interfaces diagnostics optics** command for 10-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

Table 119: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Measured laser bias current in mA.
Laser output power	Measured laser output power in mW.
Module temperature	Internally measured module temperature.
Laser rx power	Measured receive optical power in mW.
Laser bias current high alarm	Laser bias current high alarm: **On** or **Off**. Alarm ranges are vendor-specific.

Table 119: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (continued)

Field Name	Field Description
Laser bias current low alarm	Laser bias current low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power high alarm	Laser output power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power low alarm	Laser output power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp high alarm	Module temperature high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp low alarm	Module temperature low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power high alarm	Laser receive power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power low alarm	Laser receive power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser bias current high warning	Laser bias current high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current low warning	Laser bias current low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power high warning	Laser output power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power low warning	Laser output power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature high warning	Module temperature high warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature low warning	Module temperature low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power high warning	Laser receive power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power low warning	Laser receive power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current high alarm threshold	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current low alarm threshold	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.

Table 119: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (continued)

Field Name	Field Description
Laser output power high alarm threshold	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power low alarm threshold	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature high alarm threshold	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature low alarm threshold	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power high alarm threshold	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power low alarm threshold	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current high warning threshold	Laser bias current high warning threshold. Warning ranges are vendor-specific.
Laser bias current low warning threshold	Laser bias current low warning threshold. Warning ranges are vendor-specific.
Laser output power high warning threshold	Laser output power high warning threshold. Warning ranges are vendor-specific.
Laser output power low warning threshold	Laser output power low warning threshold. Warning ranges are vendor-specific.
Module temperature high warning threshold	Module temperature high warning threshold. Warning ranges are vendor-specific.
Module temperature low warning threshold	Module temperature low warning threshold. Warning ranges are vendor-specific.
Laser rx power high warning threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
Laser rx power low warning threshold	Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.

Table 120 on page 1814 lists the output fields for the **show interfaces diagnostics optics** command for XFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 120: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the **LsPOWMON** pin in hardware.
Module temperature	Temperature of the XFP optics module, in Celsius and Fahrenheit.
Laser rx power	Laser received optical power, in mW and dBm.
Laser bias current high alarm	Laser bias power setting high alarm. Displays **on** or **off**.
Laser bias current low alarm	Laser bias power setting low alarm. Displays **on** or **off**.
Laser bias current high warning	Laser bias power setting high warning. Displays **on** or **off**.
Laser bias current low warning	Laser bias power setting low warning. Displays **on** or **off**.
Laser output power high alarm	Laser output power high alarm. Displays **on** or **off**.
Laser output power low alarm	Laser output power low alarm. Displays **on** or **off**.
Laser output power high warning	Laser output power high warning. Displays **on** or **off**.
Laser output power low warning	Laser output power low warning. Displays **on** or **off**.
Module temperature high alarm	Module temperature high alarm. Displays **on** or **off**.
Module temperature low alarm	Module temperature low alarm. Displays **on** or **off**.
Module temperature high warning	Module temperature high warning. Displays **on** or **off**.
Module temperature low warning	Module temperature low warning. Displays **on** or **off**.

Table 120: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers (continued)

Field Name	Field Description
Laser rx power high alarm	Receive laser power high alarm. Displays **on** or **off**.
Laser rx power low alarm	Receive laser power low alarm. Displays **on** or **off**.
Laser rx power high warning	Receive laser power high warning. Displays **on** or **off**.
Laser rx power low warning	Receive laser power low warning. Displays **on** or **off**.
Module not ready alarm	Module not ready alarm. When **on**, indicates the module has an operational fault. Displays **on** or **off**.
Module power down alarm	Module power down alarm. When **on**, module is in a limited power mode, low for normal operation. Displays on or off.
Tx data not ready alarm	Any condition leading to invalid data on the transmit path. Displays **on** or **off**.
Tx not ready alarm	Any condition leading to invalid data on the transmit path. Displays **on** or **off**.
Tx laser fault alarm	Laser fault condition. Displays **on** or **off**.
Tx CDR loss of lock alarm	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays **on** or **off**.
Rx not ready alarm	Any condition leading to invalid data on the receive path. Displays **on** or **off**.
Rx loss of signal alarm	Receive Loss of Signal alarm. When **on**, indicates insufficient optical input power to the module. Displays **on** or **off**.
Rx CDR loss of lock alarm	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays **on** or **off**.
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm: **130.000 mA.**
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm: **10.000 mA.**
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning: **120.000 mA.**
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning: **12.000 mA.**
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm: **0.8910 mW** or **-0.50 dBm.**

Table 120: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers (continued)

Field Name	Field Description
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm: **0.2230 mW** or **-6.52 dBm.**
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning: **0.7940 mW** or **-100 dBm.**
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning: **0.2510 mW** or **-600 dBm.**
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm: **90**° C or **194**° F.
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm: **-5**° C or **23**° F.
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning: **85** ° C or **185** ° F.
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning: **0**° C or **32**° F.
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm: **1.2589 mW** or **1.00 dBm.**
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm: **0.0323 mW** or **-14.91 dBm.**
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning: **1.1220 mW** or **0.50 dBm.**
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning: **0.0363 mW** or **-14.40 dBm.**

Table 121 on page 1816 lists the output fields for the **show interfaces diagnostics optics** command for VCP. Output fields are listed in the approximate order in which they appear.

Table 121: show interfaces diagnostics optics Output for Virtual Chassis Ports

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).
Module temperature	Temperature of the optics module, in Celsius and Fahrenheit.

Table 121: show interfaces diagnostics optics Output for Virtual Chassis Ports (continued)

Field Name	Field Description
Module voltage	Internally measured module voltage.
Receiver signal average optical power	Average received optical power, in mW and dBm.
Laser bias current high alarm	Laser bias power setting high alarm. Displays **on** or **off**.
Laser bias current low alarm	Laser bias power setting low alarm. Displays **on** or **off**.
Laser bias current high warning	Laser bias power setting high warning. Displays **on** or **off**.
Laser bias current low warning	Laser bias power setting low warning. Displays **on** or **off**.
Laser output power high alarm	Laser output power high alarm. Displays **on** or **off**.
Laser output power low alarm	Laser output power low alarm. Displays **on** or **off**.
Laser output power high warning	Laser output power high warning. Displays **on** or **off**.
Laser output power low warning	Laser output power low warning. Displays **on** or **off**.
Module temperature high alarm	Module temperature high alarm. Displays **on** or **off**.
Module temperature low alarm	Module temperature low alarm. Displays **on** or **off**.
Module temperature high warning	Module temperature high warning. Displays **on** or **off**.
Module temperature low warning	Module temperature low warning. Displays **on** or **off**.
Module voltage high alarm	Module voltage high alarm. Displays **on** or **off**.
Module voltage low alarm	Module voltage low alarm. Displays **on** or **off**.
Module voltage high warning	Module voltage high warning. Displays **on** or **off**.

Table 121: show interfaces diagnostics optics Output for Virtual Chassis Ports (continued)

Field Name	Field Description
Module voltage low warning	Module voltage high warning. Displays **on** or **off**.
Laser rx power high alarm	Receive laser power high alarm. Displays **on** or **off**.
Laser rx power low alarm	Receive laser power low alarm. Displays **on** or **off**.
Laser rx power high warning	Receive laser power high warning. Displays **on** or **off**.
Laser rx power low warning	Receive laser power low warning. Displays **on** or **off**.
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning.

Table 121: show interfaces diagnostics optics Output for Virtual Chassis Ports (continued)

Field Name	Field Description
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Module voltage high alarm threshold.
Module voltage low alarm threshold	Module voltage low alarm threshold.
Module voltage high warning threshold	Module voltage high warning threshold.
Module voltage low warning threshold	Module voltage low warning threshold.
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm.
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm.
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning.
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning.

Sample Output

show interfaces diagnostics optics (DWDM and DWDM OTN)

```
user@host> show interfaces diagnostics optics ge-5/0/0
Physical interface: ge-5/0/0
  Laser bias current                      :  79.938 mA
  Laser output power                      :  1.592 mW / 2.02 dBm
  Receiver signal average optical power :  1.3854 mW / 1.42 dBm
  Laser end-of-life alarm                 :  Off
  Laser wavelength alarm                  :  Off
  Laser bias current alarm                :  Off
  Laser temperature alarm                 :  Off
  Laser power alarm                       :  Off
  Modulator temperature alarm             :  Off
  Modulator bias alarm                    :  Off
  Tx multiplexer FIFO error alarm         :  Off
  Tx loss of PLL lock alarm               :  Off
  Rx loss of average optical power alarm:  Off
  Rx loss of AC power alarm               :  Off
  Rx loss of PLL lock alarm               :  Off
```

show interfaces diagnostics optics (MPC6E with OTN MIC)

```
user@host> show interfaces diagnostics optics xe-3/0/0
Physical interface: xe-3/0/0
    Laser bias current                          :  7.806 mA
    Laser output power                          :  0.5660 mW / -2.47 dBm
    Module temperature                          :  32 degrees C / 89 degrees F
    Module voltage                              :  3.3560 V
    Receiver signal average optical power       :  0.5501 mW / -2.60 dBm
    Laser bias current high alarm               :  Off
    Laser bias current low alarm                :  Off
    Laser bias current high warning             :  Off
    Laser bias current low warning              :  Off
    Laser output power high alarm               :  Off
    Laser output power low alarm                :  Off
    Laser output power high warning             :  Off
    Laser output power low warning              :  Off
    Module temperature high alarm               :  Off
    Module temperature low alarm                :  Off
    Module temperature high warning             :  Off
    Module temperature low warning              :  Off
    Module voltage high alarm                   :  Off
    Module voltage low alarm                    :  Off
    Module voltage high warning                 :  Off
    Module voltage low warning                  :  Off
    Laser rx power high alarm                   :  Off
    Laser rx power low alarm                    :  Off
    Laser rx power high warning                 :  Off
    Laser rx power low warning                  :  Off
    Laser bias current high alarm threshold     :  11.800 mA
    Laser bias current low alarm threshold      :  4.000 mA
    Laser bias current high warning threshold   :  10.800 mA
    Laser bias current low warning threshold    :  5.000 mA
    Laser output power high alarm threshold     :  0.8310 mW / -0.80 dBm
    Laser output power low alarm threshold      :  0.2510 mW / -6.00 dBm
    Laser output power high warning threshold   :  0.6600 mW / -1.80 dBm
    Laser output power low warning threshold    :  0.3160 mW / -5.00 dBm
    Module temperature high alarm threshold     :  78 degrees C / 172 degrees F
    Module temperature low alarm threshold      :  -13 degrees C / 9 degrees F
    Module temperature high warning threshold   :  73 degrees C / 163 degrees F
    Module temperature low warning threshold    :  -8 degrees C / 18 degrees F
    Module voltage high alarm threshold         :  3.700 V
    Module voltage low alarm threshold          :  2.900 V
    Module voltage high warning threshold       :  3.600 V
    Module voltage low warning threshold        :  3.000 V
    Laser rx power high alarm threshold         :  1.0000 mW / 0.00 dBm
    Laser rx power low alarm threshold          :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold       :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold        :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (Bidirectional SFP)

```
user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
    Laser bias current                          :  13.356 mA
    Laser output power                          :  0.2210 mW / -6.56 dBm
    Module temperature                          :  36 degrees C / 96 degrees F
    Module voltage                              :  3.2180 V
    Receiver signal average optical power       :  0.2429 mW / -6.15 dBm
```



```
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  70.000 mA
    Laser bias current low alarm threshold    :  0.002 mA
    Laser bias current high warning threshold :  65.000 mA
    Laser bias current low warning threshold  :  0.002 mA
    Laser output power high alarm threshold   :  1.0000 mW / 0.00 dBm
    Laser output power low alarm threshold    :  0.0560 mW / -12.52 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0890 mW / -10.51 dBm
    Module temperature high alarm threshold   :  100 degrees C / 212 degrees F
    Module temperature low alarm threshold    :  -50 degrees C / -58 degrees F
    Module temperature high warning threshold :  95 degrees C / 203 degrees F
    Module temperature low warning threshold  :  -48 degrees C / -54 degrees F
    Module voltage high alarm threshold       :  3.700 V
    Module voltage low alarm threshold        :  2.900 V
    Module voltage high warning threshold     :  3.600 V
    Module voltage low warning threshold      :  3.000 V
    Laser rx power high alarm threshold       :  1.9953 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0001 mW / -40.00 dBm
    Laser rx power high warning threshold     :  1.0000 mW / 0.00 dBm
    Laser rx power low warning threshold      :  0.0010 mW / -30.00 dBm
```

show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-0/3/0
Physical interface: ge-0/3/0
    Laser bias current                        :  23.408 mA
    Laser output power                        :  1.479 mW / 1.70 dBm
    Module temperature                        :  37 degrees C / 99 degrees F
    Laser rx power                            :  0.121 mW / -9.16 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
```



```
Module temperature high warning          :  Off
Module temperature low warning           :  Off
Laser rx power high warning              :  Off
Laser rx power low warning               :  Off
Laser bias current high alarm threshold  :  31.000 mA
Laser bias current low alarm threshold   :  10.000 mA
Laser output power high alarm threshold  :  6.000 mW / 7.78 dBm
Laser output power low alarm threshold   :  0.100 mW / -10.00 dBm
Module temperature high alarm threshold  :  85 degrees C / 185 degrees F
Module temperature low alarm threshold   :  0 degrees C / 32 degrees F
Laser rx power high alarm threshold      :  1.000 mW / 0.00 dBm
Laser rx power low alarm threshold       :  0.001 mW / -30.00 dBm
Laser bias current high warning threshold :  28.000 mA
Laser bias current low warning threshold  :  11.000 mA
Laser output power high warning threshold :  5.000 mW / 6.99 dBm
Laser output power low warning threshold  :  0.500 mW / -3.01 dBm
Module temperature high warning threshold :  70 degrees C / 158 degrees F
Module temperature low warning threshold  :  10 degrees C / 50 degrees F
Laser rx power high warning threshold     :  0.501 mW / -3.00 dBm
Laser rx power low warning threshold      :  0.001 mW / -28.86 dBm
```

show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-1/0/0
Physical interface: ge-1/0/0
        Laser bias current                        :  49.010 mA
        Laser output power                        :  1.263 mW / 1.01 dBm
        Module temperature                        :  17 degrees C / 62 degrees F

        Module voltage                            :  4.21 V
        Laser rx power                            :  0.060 mW / -12.21 dBm
        Laser bias current high alarm             :  Off
        Laser bias current low alarm              :  Off
        Laser output power high alarm             :  Off
        Laser output power low alarm              :  Off
        Module temperature high alarm             :  Off
        Module temperature low alarm              :  Off
        Module voltage high alarm                 :  Off
        Module voltage low alarm                  :  Off
        Laser rx power high alarm                 :  Off
        Laser rx power low alarm                  :  Off
        Laser bias current high warning           :  Off
        Laser bias current low warning            :  Off
        Laser output power high warning           :  Off
        Laser output power low warning            :  Off
        Module temperature high warning           :  Off
        Module temperature low warning            :  Off
        Module voltage high warning               :  Off
        Module voltage low warning                :  Off
        Laser rx power high warning               :  Off
        Laser rx power low warning                :  Off
        Laser bias current high alarm threshold   :  70.000 mA
        Laser bias current low alarm threshold    :  20.000 mA
        Laser bias current high warning threshold :  65.000 mA
        Laser bias current low warning threshold  :  25.000 mA
        Laser output power high alarm threshold   :  1.4120 mW / 1.50 dBm
        Laser output power low alarm threshold    :  0.1990 mW / -7.01 dBm
        Laser output power high warning threshold :  1.2580 mW / 1.00 dBm
        Laser output power low warning threshold  :  0.2230 mW / -6.52 dBm
        Module temperature high alarm threshold   :  78 degrees C /172 degrees F
```

```
Module temperature low alarm threshold    :  13 degrees C / 9 degrees F
Module temperature high warning threshold :  75 degrees C /167 degrees F

Module temperature low warning threshold  :  10 degrees C / 14 degrees F

Module voltage high alarm threshold       :  5.71 V
Module voltage low alarm threshold        :  2.05 V
Module voltage high warning threshold     :  5.20 V
Module voltage low warning threshold      :  3.11 V
Laser rx power high alarm threshold       :  1.7783 mW / 2.50 dBm
Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     :  1.5849 mW / 2.00 dBm
Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (XFP and CFP Optics)

```
user@host> show interfaces diagnostics optics xe-2/1/0
Physical interface: xe-2/1/0
    Laser bias current                        :  52.060 mA
    Laser output power                        :  0.5640 mW / -2.49 dBm
    Module temperature                        :  31 degrees C / 88 degrees F
    Laser rx power                            :  0.0844 mW / -10.74 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Module not ready alarm                    :  Off
    Module power down alarm                   :  Off
    Tx data not ready alarm                   :  Off
    Tx not ready alarm                        :  Off
    Tx laser fault alarm                      :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx not ready alarm                        :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    Laser bias current high alarm threshold   :  130.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser bias current high warning threshold :  120.000 mA
    Laser bias current low warning threshold  :  12.000 mA
    Laser output power high alarm threshold   :  0.8910 mW / -0.50 dBm
    Laser output power low alarm threshold    :  0.2230 mW / -6.52 dBm
    Laser output power high warning threshold :  0.7940 mW / -1.00 dBm
    Laser output power low warning threshold  :  0.2510 mW / -6.00 dBm
    Module temperature high alarm threshold   :  90 degrees C / 194 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  85 degrees C / 185 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Laser rx power high alarm threshold       :  1.2589 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0323 mW / -14.91 dBm
```



```
    Laser rx power high warning threshold     :  1.1220 mW / 0.50 dBm
    Laser rx power low warning threshold      :  0.0363 mW / -14.40 dBm
```

show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP)

```
user@host> show interfaces diagnostics optics et-2/0/23
Physical interface: et-2/0/23
    Laser bias current                        :  8.482 mA
    Laser output power                        :  0.5890 mW / -2.30 dBm
    Module temperature                        :  51 degrees C / 123 degrees F
    Module voltage                            :  3.2970 V
    Receiver signal average optical power     :  0.5574 mW / -2.54 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  11.800 mA
    Laser bias current low alarm threshold    :  4.000 mA
    Laser bias current high warning threshold :  10.800 mA
    Laser bias current low warning threshold  :  5.000 mA
    Laser output power high alarm threshold   :  0.8310 mW / -0.80 dBm
    Laser output power low alarm threshold    :  0.2510 mW / -6.00 dBm
    Laser output power high warning threshold :  0.6600 mW / -1.80 dBm
    Laser output power low warning threshold  :  0.3160 mW / -5.00 dBm
    Module temperature high alarm threshold   :  93 degrees C / 199 degrees F
    Module temperature low alarm threshold    :  -13 degrees C / 9 degrees F
    Module temperature high warning threshold :  88 degrees C / 190 degrees F
    Module temperature low warning threshold  :  -8 degrees C / 18 degrees F
    Module voltage high alarm threshold       :  3.700 V
    Module voltage low alarm threshold        :  2.900 V
    Module voltage high warning threshold     :  3.600 V
    Module voltage low warning threshold      :  3.000 V
    Laser rx power high alarm threshold       :  1.0000 mW / 0.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics for 40-Gigabit Ethernet

```
user@host> show interfaces diagnostics optics et-7/1/0
Physical interface: et-7/1/0
    Module temperature                        :  34 degrees C / 94 degrees F
    Module voltage                            :  3.4720 V
```

```
    Module temperature high alarm            :  Off
    Module temperature low alarm             :  Off
    Module temperature high warning          :  Off
    Module temperature low warning           :  Off
    Module voltage high alarm                :  Off
    Module voltage low alarm                 :  Off
    Module voltage high warning              :  Off
    Module voltage low warning               :  Off
    Module not ready alarm                   :  Off
    Module low power alarm                   :  Off
    Module initialization incomplete alarm   :  Off
    Module fault alarm                       :  Off
    PLD Flash initialization fault alarm     :  Off
    Power supply fault alarm                 :  Off
    Checksum fault alarm                     :  Off
    Tx laser disabled alarm                  :  Off
    Tx loss of signal functionality alarm    :  Off
    Tx CDR loss of lock alarm                :  Off
    Rx loss of signal alarm                  :  Off
    Rx CDR loss of lock alarm                :  Off
    Module temperature high alarm threshold   :  80 degrees C / 176 degrees F
    Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  75 degrees C / 167 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold       :  3.5990 V
    Module voltage low alarm threshold        :  3.0000 V
    Module voltage high warning threshold     :  3.5000 V
    Module voltage low warning threshold      :  3.0990 V
    Laser bias current high alarm threshold   :  100.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser bias current high warning threshold :  80.000 mA
    Laser bias current low warning threshold  :  15.000 mA
    Laser output power high alarm threshold   :  2.8180 mW / 4.50 dBm
    Laser output power low alarm threshold    :  0.2390 mW / -6.22 dBm
    Laser output power high warning threshold :  2.2380 mW / 3.50 dBm
    Laser output power low warning threshold  :  0.3010 mW / -5.21 dBm
    Laser rx power high alarm threshold       :  2.5119 mW / 4.00 dBm
    Laser rx power low alarm threshold        :  0.0316 mW / -15.00 dBm
    Laser rx power high warning threshold     :  1.9953 mW / 3.00 dBm
    Laser rx power low warning threshold      :  0.0631 mW / -12.00 dBm
    Laser temperature high alarm threshold    :  80 degrees C / 176 degrees F
    Laser temperature low alarm threshold     :  -10 degrees C / 14 degrees F
    Laser temperature high warning threshold  :  75 degrees C / 167 degrees F
    Laser temperature low warning threshold   :  -5 degrees C / 23 degrees F
  Lane 0
    Laser bias current                        :  27.829 mA
    Laser output power                        :  0.851 mW / -0.70 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.894 mW / -0.49 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
```



```
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Tx CDR loss of lock alarm                :  Off
    Rx loss of signal alarm                  :  Off
    Rx CDR loss of lock alarm                :  Off
    APD supply fault alarm                   :  Off
    TEC fault alarm                          :  Off
    Wavelength unlocked alarm                :  Off
  Lane 1
    Laser bias current                       :  35.374 mA
    Laser output power                       :  0.896 mW / -0.48 dBm
    Laser temperature                        :  34 degrees C / 94 degrees F
    Laser receiver power                     :  0.707 mW / -1.50 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser output power high alarm            :  Off
    Laser output power low alarm             :  Off
    Laser output power high warning          :  Off
    Laser output power low warning           :  Off
    Laser temperature high alarm             :  Off
    Laser temperature low alarm              :  Off
    Laser temperature high warning           :  Off
    Laser temperature low warning            :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
    Tx CDR loss of lock alarm                :  Off
    Rx loss of signal alarm                  :  Off
    Rx CDR loss of lock alarm                :  Off
    APD supply fault alarm                   :  Off
    TEC fault alarm                          :  Off
    Wavelength unlocked alarm                :  Off
  Lane 2
    Laser bias current                       :  29.173 mA
    Laser output power                       :  0.890 mW / -0.51 dBm
    Laser temperature                        :  34 degrees C / 94 degrees F
    Laser receiver power                     :  0.704 mW / -1.52 dBm
    Laser bias current high alarm            :  Off
    Laser bias current low alarm             :  Off
    Laser bias current high warning          :  Off
    Laser bias current low warning           :  Off
    Laser output power high alarm            :  Off
    Laser output power low alarm             :  Off
    Laser output power high warning          :  Off
    Laser output power low warning           :  Off
    Laser temperature high alarm             :  Off
    Laser temperature low alarm              :  Off
    Laser temperature high warning           :  Off
    Laser temperature low warning            :  Off
    Laser receiver power high alarm          :  Off
    Laser receiver power low alarm           :  Off
    Laser receiver power high warning        :  Off
    Laser receiver power low warning         :  Off
    Tx loss of signal functionality alarm    :  Off
```



```
      Tx CDR loss of lock alarm                :  Off
      Rx loss of signal alarm                  :  Off
      Rx CDR loss of lock alarm                :  Off
      APD supply fault alarm                   :  Off
      TEC fault alarm                          :  Off
      Wavelength unlocked alarm                :  Off
    Lane 3
      Laser bias current                       :  36.164 mA
      Laser output power                       :  0.899 mW / -0.46 dBm
      Laser temperature                        :  34 degrees C / 94 degrees F
      Laser receiver power                     :  0.892 mW / -0.50 dBm
      Laser bias current high alarm            :  Off
      Laser bias current low alarm             :  Off
      Laser bias current high warning          :  Off
      Laser bias current low warning           :  Off
      Laser output power high alarm            :  Off
      Laser output power low alarm             :  Off
      Laser output power high warning          :  Off
      Laser output power low warning           :  Off
      Laser temperature high alarm             :  Off
      Laser temperature low alarm              :  Off
      Laser temperature high warning           :  Off
      Laser temperature low warning            :  Off
      Laser receiver power high alarm          :  Off
      Laser receiver power low alarm           :  Off
      Laser receiver power high warning        :  Off
      Laser receiver power low warning         :  Off
      Tx loss of signal functionality alarm    :  Off
      Tx CDR loss of lock alarm                :  Off
      Rx loss of signal alarm                  :  Off
      Rx CDR loss of lock alarm                :  Off
      APD supply fault alarm                   :  Off
      TEC fault alarm                          :  Off
      Wavelength unlocked alarm                :  Off
```

show interfaces diagnostics optics (P1-PTX-2-100G-WDM)

```
user@host> show interfaces diagnostics optics et-1/0/0
Physical interface: et-1/0/0
    Module temperature                         :  37 degrees C / 98 degrees F
    Module voltage                             :  3.3370 V
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Module temperature high warning            :  Off
    Module temperature low warning             :  Off
    Module voltage high alarm                  :  Off
    Module voltage low alarm                   :  Off
    Module voltage high warning                :  Off
    Module voltage low warning                 :  Off
    Module not ready alarm                     :  Off
    Module low power alarm                     :  Off
    Module initialization incomplete alarm     :  Off
    Module fault alarm                         :  Off
    PLD Flash initialization fault alarm       :  Off
    Power supply fault alarm                   :  Off
    Checksum fault alarm                       :  Off
    Tx laser disabled alarm                    :  Off
    Tx loss of signal functionality alarm      :  Off
    Tx CDR loss of lock alarm                  :  Off
    Rx loss of signal alarm                    :  Off
    Rx CDR loss of lock alarm                  :  Off
```

```
    Module temperature high alarm threshold   :  70 degrees C / 158 degrees F
    Module temperature low alarm threshold    :  0 degrees C / 32 degrees F
    Module temperature high warning threshold :  68 degrees C / 154 degrees F
    Module temperature low warning threshold  :  2 degrees C / 36 degrees F
    Module voltage high alarm threshold       :  3.4640 V
Module voltage low alarm threshold        :  3.1340 V
    Module voltage high warning threshold     :  3.4310 V
    Module voltage low warning threshold      :  3.1670 V
    Laser bias current high alarm threshold   :  300.000 mA
    Laser bias current low alarm threshold    :  75.000 mA
    Laser bias current high warning threshold :  287.500 mA
    Laser bias current low warning threshold  :  87.500 mA
Rx power high alarm threshold       :  2.8184 mW / 4.50 dBm
    Rx power low alarm threshold        :  0.0251 mW / -16.00 dBm
    Rx power high warning threshold     :  2.5119 mW / 4.00 dBm
    Rx power low warning threshold      :  0.0501 mW / -13.00 dBm
    LOS alarm threshold             : 0.0158mW/ -18.01 dBm
    LOS warning threshold            : 0.0251mW/ -16.00 DBm
    Laser temperature high alarm threshold    :  57 degrees C / 135 degrees F
    Laser temperature low alarm threshold     :  25 degrees C / 77 degrees F
    Laser temperature high warning threshold  :  55 degrees C / 131 degrees F
    Laser temperature low warning threshold   :  27 degrees C / 81 degrees F
  Lane 0
    Laser bias current                        :  164.384 mA
Tx power                     :  1.181 mW / 0.72 dBm
    Laser temperature                         :  41 degrees C / 106 degrees F
    Rx power                    :  0.632 mW / -1.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Tx power high alarm               :  Off
    Tx power low alarm                :  Off
    Tx power high warning             :  Off
    Tx power low warning              :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Rx  power high alarm                :  Off
    Rx power low alarm                 :  Off
    Rx power high warning             :  Off
    Rx power low warning              :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
```

show interfaces diagnostics optics (P1-PTX-24-10G-W-SFPP)

```
user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
    Laser bias current                        :  13.356 mA
    Laser output power                        :  0.2210 mW / -6.56 dBm
    Module temperature                        :  36 degrees C / 96 degrees F
    Module voltage                            :  3.2180 V
    Receiver signal average optical power     :  0.2429 mW / -6.15 dBm
    Wavelength Channel number                 :  1
```



```
    Wavelength setpoint                        :  1568.80 nm
    Tx Dither                                  :  Disabled
    Frequency Error                            :  0.00 GHz
    Wavelength Error                           :  0.00 nm
    Laser bias current high alarm              :  Off
    Laser bias current low alarm               :  Off
    Laser bias current high warning            :  Off
    Laser bias current low warning             :  Off
    Laser output power high alarm              :  Off
    Laser output power low alarm               :  Off
    Laser output power high warning            :  Off
    Laser output power low warning             :  Off
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Module temperature high warning            :  Off
    Module temperature low warning             :  Off
    Module voltage high alarm                  :  Off
    Module voltage low alarm                   :  Off
    Module voltage high warning                :  Off
    Module voltage low warning                 :  Off
    Laser rx power high alarm                  :  Off
    Laser rx power low alarm                   :  Off
    Laser rx power high warning                :  Off
    Laser rx power low warning                 :  Off
    TEC fault alarm                            :  Off
    Wavelength unlocked alarm                  :  Off
    Tx Tune                                    :  Off
    Laser bias current high alarm threshold    :  70.000 mA
    Laser bias current low alarm threshold     :  0.002 mA
    Laser bias current high warning threshold :  65.000 mA
    Laser bias current low warning threshold   :  0.002 mA
    Laser output power high alarm threshold    :  1.0000 mW / 0.00 dBm
    Laser output power low alarm threshold     :  0.0560 mW / -12.52 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold   :  0.0890 mW / -10.51 dBm
    Module temperature high alarm threshold    :  100 degrees C / 212 degrees F
    Module temperature low alarm threshold     :  -50 degrees C / -58 degrees F
    Module temperature high warning threshold :  95 degrees C / 203 degrees F
    Module temperature low warning threshold   :  -48 degrees C / -54 degrees F
    Module voltage high alarm threshold        :  3.700 V
    Module voltage low alarm threshold         :  2.900 V
    Module voltage high warning threshold      :  3.600 V
    Module voltage low warning threshold       :  3.000 V
    Laser rx power high alarm threshold        :  1.9953 mW / 3.00 dBm
    Laser rx power low alarm threshold         :  0.0001 mW / -40.00 dBm
    Laser rx power high warning threshold      :  1.0000 mW / 0.00 dBm
    Laser rx power low warning threshold       :  0.0010 mW / -30.00 dBm
```

show interfaces diagnostics optics (P2-10G-40G-QSFPP PIC in 40-Gigabit Ethernet mode)

```
user@host> show interfaces diagnostics optics et-0/1/5
Physical interface: et-0/1/5
    Module temperature                         :  30 degrees C / 85 degrees F
    Module voltage                             :  3.2760 V
    Module temperature high alarm              :  Off
    Module temperature low alarm               :  Off
    Module temperature high warning            :  Off
    Module temperature low warning             :  Off
    Module voltage high alarm                  :  Off
    Module voltage low alarm                   :  Off
    Module voltage high warning                :  Off
```

```
    Module voltage low warning               :  Off
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  5 degrees C / 41 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  3.4640 V
    Module voltage low warning threshold      :  3.1340 V
    Laser bias current high alarm threshold   :  10.000 mA
    Laser bias current low alarm threshold    :  0.500 mA
    Laser bias current high warning threshold :  9.500 mA
    Laser bias current low warning threshold  :  1.000 mA
    Laser output power high alarm threshold   :  0.0000 mW / - Inf dBm
    Laser output power low alarm threshold    :  0.0000 mW / - Inf dBm
    Laser output power high warning threshold :  0.0000 mW / - Inf dBm
    Laser output power low warning threshold  :  0.0000 mW / - Inf dBm
    Laser rx power high alarm threshold       :  2.1878 mW / 3.40 dBm
    Laser rx power low alarm threshold        :  0.0446 mW / -13.51 dBm
    Laser rx power high warning threshold     :  1.7378 mW / 2.40 dBm
    Laser rx power low warning threshold      :  0.1122 mW / -9.50 dBm
  Lane 0
    Laser bias current                        :  7.065 mA
    Laser output power                        :  0.710 mW / -1.49 dBm
    Laser receiver power                      :  0.472 mW / -3.26 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
  Lane 1
    Laser bias current                        :  6.978 mA
    Laser output power                        :  0.771 mW / -1.13 dBm
    Laser receiver power                      :  0.450 mW / -3.47 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
  Lane 2
    Laser bias current                        :  6.955 mA
    Laser output power                        :  0.760 mW / -1.19 dBm
    Laser receiver power                      :  0.556 mW / -2.55 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
```



```
      Tx loss of signal functionality alarm     :  Off
      Rx loss of signal alarm                   :  Off
    Lane 3
      Laser bias current                        :  6.981 mA
      Laser output power                        :  0.736 mW / -1.33 dBm
      Laser receiver power                      :  0.537 mW / -2.70 dBm
      Laser bias current high alarm             :  Off
      Laser bias current low alarm              :  Off
      Laser bias current high warning           :  Off
      Laser bias current low warning            :  Off
      Laser receiver power high alarm           :  Off
      Laser receiver power low alarm            :  Off
      Laser receiver power high warning         :  Off
      Laser receiver power low warning          :  Off
      Tx loss of signal functionality alarm     :  Off
      Rx loss of signal alarm                   :  Off
```

show interfaces diagnostics optics (P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode)

```
user@host> show interfaces diagnostics optics et-0/1/5:3
Physical interface: et-0/1/5:3
      Module temperature                        :  30 degrees C / 85 degrees F
      Module voltage                            :  3.2760 V
      Module temperature high alarm             :  Off
      Module temperature low alarm              :  Off
      Module temperature high warning           :  Off
      Module temperature low warning            :  Off
      Module voltage high alarm                 :  Off
      Module voltage low alarm                  :  Off
      Module voltage high warning               :  Off
      Module voltage low warning                :  Off
      Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
      Module temperature low alarm threshold    :  5 degrees C / 41 degrees F
      Module temperature high warning threshold :  70 degrees C / 158 degrees F
      Module temperature low warning threshold  :  0 degrees C / 32 degrees F
      Module voltage high alarm threshold       :  3.6300 V
      Module voltage low alarm threshold        :  2.9700 V
      Module voltage high warning threshold     :  3.4640 V
      Module voltage low warning threshold      :  3.1340 V
      Laser bias current high alarm threshold   :  10.000 mA
      Laser bias current low alarm threshold    :  0.500 mA
      Laser bias current high warning threshold :  9.500 mA
      Laser bias current low warning threshold  :  1.000 mA
      Laser output power high alarm threshold   :  0.0000 mW / - Inf dBm
      Laser output power low alarm threshold    :  0.0000 mW / - Inf dBm
      Laser output power high warning threshold :  0.0000 mW / - Inf dBm
      Laser output power low warning threshold  :  0.0000 mW / - Inf dBm
      Laser rx power high alarm threshold       :  2.1878 mW / 3.40 dBm
      Laser rx power low alarm threshold        :  0.0446 mW / -13.51 dBm
      Laser rx power high warning threshold     :  1.7378 mW / 2.40 dBm
      Laser rx power low warning threshold      :  0.1122 mW / -9.50 dBm
    Lane 3
      Laser bias current                        :  6.981 mA
      Laser output power                        :  0.736 mW / -1.33 dBm
      Laser receiver power                      :  0.537 mW / -2.70 dBm
      Laser bias current high alarm             :  Off
      Laser bias current low alarm              :  Off
      Laser bias current high warning           :  Off
      Laser bias current low warning            :  Off
      Laser receiver power high alarm           :  Off
      Laser receiver power low alarm            :  Off
```

```
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
```

show interfaces diagnostics optics (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host> show interfaces diagnostics optics et-2/0/0
Physical interface: et-2/0/0
    Module temperature                        :  39 degrees C / 102 degrees F
    Module voltage                            :  3.2300 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module not ready alarm                    :  Off
    Module low power alarm                    :  Off
    Module initialization incomplete alarm    :  Off
    Module fault alarm                        :  Off
    Tx laser disabled alarm                   :  Off
    Rx loss of signal alarm                   :  Off
    Module temperature high alarm threshold   :  90 degrees C / 194 degrees F
    Module temperature low alarm threshold    :  -20 degrees C / -4 degrees F
    Module temperature high warning threshold :  0 degrees C / 32 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  0.0000 V
    Module voltage low warning threshold      :  0.0000 V
    Rx power high alarm threshold             :  6.5535 mW / 8.16 dBm
    Rx power low alarm threshold              :  0.0028 mW / -25.53 dBm
    Rx power high warning threshold           :  6.5535 mW / 8.16 dBm
    Rx power low warning threshold            :  0.0028 mW / -25.53 dBm
    LOS alarm threshold                       :  0.0028 mW / -25.53 dBm
    LOS warning threshold                     :  0.0028 mW / -25.53 dBm
    Modem lock state                          :  OK
  Lane 0
    Tx power                                  :  1.000 mW / 0.00 dBm
    Module temperature                        :  51 degrees C / 124 degrees F
    Rx power (total)                          :  0.644 mW / -1.91 dBm
    Rx power (signal)                         :  0.618 mW / -2.09 dBm
    Lane chromatic dispersion                 :  -22 ps/nm
    Lane differential group delay             :  5 ps
    Lane Q2 factor                            :  14.20 dB
    Lane carrier frequency offset             :  -534 Mz
    Lane electrical SNR                       :  9.20 dB
    Tx power high alarm                       :  Off
    Tx power low alarm                        :  Off
    Tx power high warning                     :  Off
    Tx power low warning                      :  Off
    Rx power high alarm                       :  Off
    Rx power low alarm                        :  Off
    Rx power high warning                     :  Off
    Rx power low warning                      :  Off
    Rx loss of signal alarm                   :  Off
    Wavelength unlocked alarm                 :  Off
```


show interfaces diagnostics optics (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC

```
user@host> show interfaces diagnostics optics et-4/0/0
Physical interface: et-4/0/0
    Laser output power                          :  54 degrees C / 129 degrees F
    Tx module temperature                       :  0.0000
    Module temperature high alarm               :  Off
    Module temperature low alarm                :  Off
    Module temperature high warning             :  Off
    Module temperature low warning              :  Off
    Module voltage high alarm                   :  Off
    Module voltage low alarm                    :  Off
    Module voltage high warning                 :  Off
    Module voltage low warning                  :  Off
    Module not ready alarm                      :  Off
    Module low power alarm                      :  Off
    Module initialization incomplete alarm      :  Off
    Module fault alarm                          :  Off
    PLD Flash initialization fault alarm        :  Off
    Power supply fault alarm                    :  Off
    Checksum fault alarm                        :  Off
    Tx laser disabled alarm                     :  Off
    Rx loss of signal alarm                     :  Off
    Module temperature high alarm threshold     :  80 degrees C / 176 degrees F
    Module temperature low alarm threshold      :  0 degrees C / 32 degrees F
    Module temperature high warning threshold :  65 degrees C / 149 degrees F
    Module temperature low warning threshold    :  5 degrees C / 41 degrees F
    Module voltage high alarm threshold         :  0.0000 V
    Module voltage low alarm threshold          :  0.0000 V
    Module voltage high warning threshold       :  0.0000 V
    Module voltage low warning threshold        :  0.0000 V
    Rx power high alarm threshold               :  0.0000 mW / - Inf dBm
    Rx power low alarm threshold                :  0.0000 mW / - Inf dBm
    Rx power high warning threshold             :  0.0000 mW / - Inf dBm
    Rx power low warning threshold              :  0.0000 mW / - Inf dBm
    LOS alarm threshold                         :  0.0158 mW / -18.01 dBm
    LOS warning threshold                       :  0.0251 mW / -16.00 dBm
    Modem lock state                            :  OK
  Lane 0
    Tx power                                    :  1.000 mW / 0.00 dBm
    Module temperature                          :  0 degrees C / 32 degrees F
    Rx power (total)                            :  0.000 mW / - Inf dBm
    Rx power (signal)                           :  0.999 mW / -0.00 dBm
    Lane chromatic dispersion                   :  6 ps/nm
    Lane differential group delay               :  3 ps
    Lane Q2 factor                              :  15.40 dB
    Lane carrier frequency offset               :  0 MHz
    Lane electrical SNR                         :  16.60 dB
    Tx power high alarm                         :  Off
    Tx power low alarm                          :  Off
    Tx power high warning                       :  Off
    Tx power low warning                        :  Off
    Rx power high alarm                         :  Off
    Rx power low alarm                          :  Off
    Rx power high warning                       :  Off
Rx power low warning                        :  Off
    Rx loss of signal alarm                     :  Off
    Wavelength unlocked alarm                   :  Off
    Laser end-of-life alarm                     :  Off
  Lane 1
```

```
      Tx power                                : 1.000 mW / 0.00 dBm
      Module temperature                      : 0 degrees C / 32 degrees F
      Rx power (total)                        : 0.000 mW / - Inf dBm
      Rx power (signal)                       : 0.999 mW / -0.00 dBm
      Tx power high alarm                     : Off
      Tx power low alarm                      : Off
      Tx power high warning                   : Off
      Tx power low warning                    : Off
      Rx power high alarm                     : Off
      Rx power low alarm                      : Off
      Rx power high warning                   : Off
      Rx power low warning                    : Off
      Rx loss of signal alarm                 : Off
      Wavelength unlocked alarm               : Off
      Laser end-of-life alarm                 : Off
    Lane 2
      Tx power                                : 1.000 mW / 0.00 dBm
      Module temperature                      : 0 degrees C / 32 degrees F
      Rx power (total)                        : 0.000 mW / - Inf dBm
      Rx power (signal)                       : 0.999 mW / -0.00 dBm
      Tx power high alarm                     : Off
      Tx power low alarm                      : Off
      Tx power high warning                   : Off
      Tx power low warning                    : Off
      Rx power high alarm                     : Off
      Rx power low alarm                      : Off
      Rx power high warning                   : Off
      Rx power low warning                    : Off
      Rx loss of signal alarm                 : Off
      Wavelength unlocked alarm               : Off
      Laser end-of-life alarm                 : Off
    Lane 3
      Tx power                                : 1.000 mW / 0.00 dBm
      Module temperature                      : 0 degrees C / 32 degrees F
      Rx power (total)                        : 0.000 mW / - Inf dBm
      Rx power (signal)                       : 0.999 mW / -0.00 dBm
      Tx power high alarm                     : Off
      Tx power low alarm                      : Off
      Tx power high warning                   : Off
      Tx power low warning                    : Off
  Rx power high alarm                     : Off
      Rx power low alarm                      : Off
      Rx power high warning                   : Off
      Rx power low warning                    : Off
      Rx loss of signal alarm                 : Off
      Wavelength unlocked alarm               : Off
      Laser end-of-life alarm                 : Off
```

show interfaces diagnostics optics (for VCP)

```
user@host> show interfaces diagnostics optics vcp-2/0/1
Physical interface: vcp-2/0/1
    Laser bias current                        : 5.494 mA
    Laser output power                        : 0.2960 mW / -5.29 dBm
    Module temperature                        : 22 degrees C / 71 degrees F
    Module voltage                            : 3.2810 V
    Receiver signal average optical power     : 0.2426 mW / -6.15 dBm
    Laser bias current high alarm             : Off
    Laser bias current low alarm              : Off
    Laser bias current high warning           : Off
    Laser bias current low warning            : Off
```

```
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (MPC7 with interfaces disabled)

```
user@host> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
    Module temperature                        :  34 degrees C / 93 degrees F
    Module voltage                            :  3.2660 V
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Module temperature high alarm threshold   :  75 degrees C / 167 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Module voltage high alarm threshold       :  3.6300 V
    Module voltage low alarm threshold        :  2.9700 V
    Module voltage high warning threshold     :  3.4640 V
    Module voltage low warning threshold      :  3.1340 V
    Laser bias current high alarm threshold   :  9.999 mA
```



```
    Laser bias current low alarm threshold    :  0.499 mA
    Laser bias current high warning threshold :  9.499 mA
    Laser bias current low warning threshold  :  0.999 mA
    Laser output power high alarm threshold   :  0.0000 mW / - Inf dBm
    Laser output power low alarm threshold    :  0.0000 mW / - Inf dBm
    Laser output power high warning threshold :  0.0000 mW / - Inf dBm
    Laser output power low warning threshold  :  0.0000 mW / - Inf dBm
    Laser rx power high alarm threshold       :  2.1878 mW / 3.40 dBm
    Laser rx power low alarm threshold        :  0.0446 mW / -13.51 dBm
    Laser rx power high warning threshold     :  1.7378 mW / 2.40 dBm
    Laser rx power low warning threshold      :  0.1122 mW / -9.50 dBm
  Lane 0
    Laser bias current                        :  6.697 mA
    Laser output power                        :  0.738 mW / -1.32 dBm
    Laser receiver power                      :  0.790 mW / -1.02 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
  Lane 1
    Laser bias current                        :  6.961 mA
    Laser output power                        :  0.908 mW / -0.42 dBm
    Laser receiver power                      :  0.827 mW / -0.83 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
  Lane 2
    Laser bias current                        :  6.926 mA
    Laser output power                        :  0.888 mW / -0.51 dBm
    Laser receiver power                      :  0.820 mW / -0.86 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Rx loss of signal alarm                   :  Off
    Tx laser disabled alarm                   :  Off
  Lane 3
    Laser bias current                        :  6.817 mA
    Laser output power                        :  0.846 mW / -0.73 dBm
    Laser receiver power                      :  0.827 mW / -0.82 dBm
    Laser bias current high alarm             :  Off
```



```
Laser bias current low alarm             :  Off
Laser bias current high warning          :  Off
Laser bias current low warning           :  Off
Laser receiver power high alarm          :  Off
Laser receiver power low alarm           :  Off
Laser receiver power high warning        :  Off
Laser receiver power low warning         :  Off
Tx loss of signal functionality alarm    :  Off
Rx loss of signal alarm                  :  Off
Tx laser disabled alarm                  :  Off
```

show interfaces (far-end-interval)

Syntax show interfaces far-end-interval *interface-fpc/pic/port*

Release Information Command introduced in Junos OS Release 9.4.

Description On channelized interfaces, display the far end interval data for the specified interface.

Required Privilege Level view

List of Sample Output
show interfaces far-end-interval coc12-5/2/0 on page 1838
show interfaces far-end-interval coc1-5/2/1:1 on page 1839

Output Fields Table 122 on page 1838 lists the output fields for the **show interfaces far-end-interval** command. Output fields are listed in the approximate order in which they appear.

Table 122: show interfaces far-end-interval Output Fields

Field Name	Field Description
Physical interface	Interface FPC/PIC/port values.
SNMP ifIndex	SNMP interface index value.
ES-L/P	Error detection—Errored seconds.
SES-L/P	Error detection—Severely errored seconds.
UAS-L/P	Error detection—Unavailable seconds.

Sample Output

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```


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 1853
show interfaces brief (Fast Ethernet) on page 1854
show interfaces detail (Fast Ethernet) on page 1854
show interfaces extensive (Fast Ethernet) on page 1854

Output Fields Table 123 on page 1840 lists the output fields for the **show interfaces** (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 123: 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

Table 123: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
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
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

Table 123: show interfaces Fast Ethernet Output Fields (continued)

Field Name	Field Description	Level of Output
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	(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 command.	**detail extensive**

Table 123: 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 123: 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 123: 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 123: 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 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 123: 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 123: 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) • **SES-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 123: 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 123: 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 123: 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 123: 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 123: 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

List of Syntax

Syntax (Gigabit Ethernet) on page 1857
Syntax (10 Gigabit Ethernet) on page 1857
Syntax (SRX Series Devices) on page 1857

Syntax (Gigabit Ethernet)

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

Syntax (10 Gigabit Ethernet)

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

Syntax (SRX Series Devices)

show interfaces (
<*interface-name*>
<brief | detail | extensive | terse>
<controller *interface-name*>|
<descriptions *interface-name*>|
<destination-class (all | *destination-class-name logical-interface-name*)>|
<diagnostics optics *interface-name*>|
<far-end-interval *interface-fpc/pic/port*>|
<filters *interface-name*>|
<flow-statistics *interface-name*>|
<interval *interface-name*>|
<load-balancing (detail | *interface-name*)>|
<mac-database mac-address *mac-address*>|
<mc-ae id *identifier* unit *number* revertive-info>|
<media *interface-name*>|
<policers *interface-name*>|
<queue both-ingress-egress egress forwarding-class *forwarding-class* ingress l2-statistics>|
<redundancy (detail | *interface-name*)>|
<routing brief detail summary *interface-name*>|
<routing-instance (all | *instance-name*)>|
<snmp-index *snmp-index*>|
<source-class (all | *destination-class-name logical-interface-name*)>|
<statistics *interface-name*>|
<switch-port *switch-port number*>|
<transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all | *interface-name*)>|
<zone *interface-name*>
)

Release Information

Command introduced before Junos OS Release 7.4 for Gigabit interfaces.
Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.
Command modified in Junos OS Release 9.5 for SRX Series devices.

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

Description Display status information about the specified Gigabit Ethernet interface.

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

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

NOTE: On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Options For Gigabit interfaces:

ge-*fpc/pic/port*—Display standard information about the specified 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.

For 10 Gigabit interfaces:

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.

For SRX interfaces:

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
  - **at-** ***pim/0/port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim/0/ port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim/0/port***—Channelized T1 interface.
  - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
  - **e1-*pim/0/port***—E1 interface.
  - **e3-*pim/0/port***—E3 interface.
  - **fe-*pim/0/port***—Fast Ethernet interface.
  - **ge-*pim/0/port***—Gigabit Ethernet interface.
  - **se-*pim/0/port***—Serial interface.
  - **t1-*pim/0/port***—T1 (also called DS1) interface.
  - **t3-*pim/0/port***—T3 (also called DS3) interface.
  - **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
  - **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
    - **at-** ***pim/0/port***—ATM-over-ADSL or ATM-over-SHDSL interface.
    - **ce1-*pim/0/ port***—Channelized E1 interface.
    - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
    - **ct1-*pim/0/port***—Channelized T1 interface.
    - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
    - **e1-*pim/0/port***—E1 interface.
    - **e3-*pim/0/port***—E3 interface.
    - **fe-*pim/0/port***—Fast Ethernet interface.
    - **ge-*pim/0/port***—Gigabit Ethernet interface.
    - **se-*pim/0/port***—Serial interface.
    - **t1-*pim/0/port***—T1 (also called DS1) interface.

- **t3-*pim*/0/*port***—T3 (also called DS3) interface.
- **wx-*slot*/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Additional Information In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level view

Related Documentation

- *Understanding Layer 2 Interfaces on Security Devices*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

List of Sample Output

Output Fields

Table 124 on page 1861 describes the output fields for the **show interfaces** (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 125 on page 1889.

Table 124: 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

Table 124: 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**
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
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

Table 124: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
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 (ms).	**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
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
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**
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 31 under the show interfaces command.	**detail extensive**

Table 124: show interfaces (Gigabit 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 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 124: 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. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Drops** field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the 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 scenario, 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**

Table 124: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
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. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Dropped packets** field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.	**detail extensive**
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 router configuration, an alarm can ring the red or yellow alarm bell on the router, 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
Interface transmit statistics	(On MX Series devices) Status of the **interface-transmit-statistics** configuration: Enabled or Disabled. • **Enabled**—When the **interface-transmit-statistics** statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • **Disabled**—When the **interface-transmit-statistics** statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics: • **Corrected Errors**—Count of corrected errors in the last second. • **Corrected Error Ratio**—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**

Table 124: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

Link Degrade	Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.
	• **Link Monitoring**—Indicates if physical link degrade monitoring is enabled on the interface.
	  ◦ **Enable**—Indicates that link degrade monitoring has been enabled (using the **link-degrade-monitor** statement) on the interface.
	  ◦ **Disable**—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds.
	• **Link Degrade Set Threshold**—The BER threshold value at which the link is considered degraded and a corrective action is triggered.
	• **Link Degrade Clear Threshold**—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted.
	• **Estimated BER**—The estimated bit error rate.
	• **Link-degrade event**—Shows link degrade event information.
	  ◦ **Seconds**—Time (in seconds) elapsed after a link degrade event occurred.
	  ◦ **Count**—The number of link degrade events recorded.
	  ◦ **State**—Shows the link degrade status (example: **Defect Active**).
		detail extensive

Table 124: 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 31 under the show interfaces 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**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the **VLAN tagged frames** field displays 0 when the **show interfaces** command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.
	• **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 124: 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 may 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 router 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 router (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must 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 must not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0.	**extensive**
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 124: 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) • **SES-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. Any 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 124: 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 (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 124: 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**—Information from the remote Ethernet device:
	  ◦ **Link mode**—Depending on the capability of the link partner, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), **Symmetric/Asymmetric** (link partner supports **PAUSE** on receive and transmit or only **PAUSE** on transmit), and **None** (link partner does not support flow control).
	  ◦ **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 local Ethernet device:
	  ◦ **Flow control**—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are **Symmetric/Asymmetric** (local device supports **PAUSE** on receive and transmit or only **PAUSE** on receive) and **None** (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display **Symmetric** (local device supports **PAUSE** on receive and transmit), **Asymmetric** (local device supports **PAUSE** on receive), and **None** (local device does not support flow control).
	  ◦ **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 router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 124: 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
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 124: show interfaces (Gigabit 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
ACI VLAN	Information displayed for agent circuit identifier (ACI) interface set configured with the **agent-circuit-id** autoconfiguration stanza. **Dynamic Profile**—Name of the dynamic profile that defines the ACI interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information. NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the **line-identity** autoconfiguration stanza.	**brief detail extensive** none
Line Identity	Information displayed for access-line-identifier (ALI) interface sets configured with the **line-identity** autoconfiguration stanza. • **Dynamic Profile**—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: **Circuit-id**, **Remote-id**, or **Accept-no-ids**. More than one option can be configured. If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information. NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the **agent-circuit-id** autoconfiguration stanza.	**detail**
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
Neighbor Discovery Protocol (NDP)Queue Statistics	NDP statistics for protocol **inet6** under logical interface statistics. • **Max nh cache**—Maximum interface neighbor discovery nexthop cache size. • **New hold nh limit**—Maximum number of new unresolved nexthops. • **Curr nh cnt**—Current number of resolved nexthops in the NDP queue. • **Curr new hold cnt**—Current number of unresolved nexthops in the NDP queue. • **NH drop cnt**—Number of NDP requests not serviced.	All levels
Dynamic Profile	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	**detail extensive** none
Service Name Table	Name of the service name table for the interface configured with a PPPoE family.	**detail extensive** none

Table 124: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	**detail extensive** none
Duplicate Protection	State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	**detail extensive** none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**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 router.	**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**

Table 124: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
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 parentheses 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 parentheses 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 the 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 interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

The following table describes the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
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
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**
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.	**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
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**
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
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**
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. • **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**

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.	**extensive**

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**

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**

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**
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
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**

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
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table 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). 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 125: Gigabit and 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 include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 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	

Table 126 on page 1890 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 126: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Current address	Configured MAC address.	**detail extensive** none

Table 126: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour*:*minute*:*second*:*timezone*** **(*hour*:*minute*:*second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	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.	**detail extensive**

Table 126: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** . • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 126: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

MAC statistics	**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.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—There are two possible conditions regarding the number of oversized frames:
	  ◦ Packet length exceeds 1518 octets, or
	  ◦ Packet length exceeds MRU
	• **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
	• **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.
	• **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.
	• **Code violations**—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."
		extensive

Table 126: show interfaces 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 device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0.	**extensive**

Autonegotiation information	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.
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 126: show interfaces 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**
Interface transmit statistics	Status of the **interface-transmit-statistics** configuration: Enabled or Disabled.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**

Table 126: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Security	Security zones that interface belongs to.	**extensive**
Flow Input statistics	Statistics on packets received by flow module.	**extensive**
Flow Output statistics	Statistics on packets sent by flow module.	**extensive**
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	**extensive**
Protocol	Protocol family.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. .	**detail extensive**
Addresses, Flags	Information about the address flags..	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output Gigabit Ethernet

show interfaces (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
 Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
        Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
   Protocol inet6, MTU: 1500
    Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
cnt: 4, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 2001:db8:/32, Local: 2001:db8::5
      Addresses, Flags: Is-Preferred
        Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces (link degrade status)

```
user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
  Interface index: 157, SNMP ifIndex: 537
  Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU 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
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
  Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade* :
  Link Monitoring                   :  Enable
  Link Degrade Set Threshold:       :  1E-7
  Link Degrade Clear Threshold:     :  1E-12
  Estimated BER                     :  1E-7
  Link-degrade event                :  Seconds     Count    State
                                           782         1     Defect Active
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
   Output bytes  :        240614363944           772721536 bps
   Output packets:          3538446506             1420444 pps
    Direction : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
     Output bytes  :        195560312716           522726272 bps
     Output packets:          4251311146             1420451 pps

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
```

```
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
   Input  bytes  :              271524
   Output bytes  :            37769598
   Input  packets:                3664
   Output packets:              885790
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :           16681118
    Input  packets:                  0
    Output packets:             362633
  Local statistics:
   Input  bytes  :              271524
   Output bytes  :              308560
   Input  packets:                3664
   Output packets:                3659
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :            37461038                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:              882131                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0     0 bps
    Output bytes  :           16681118     0 bps
    Input  packets:                  0     0 pps
    Output packets:             362633     0 pps
```

show interfaces brief (Gigabit Ethernet)

user@host> **show interfaces ge-3/0/2 brief**

```
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

show interfaces detail (Gigabit Ethernet)

user@host> **show interfaces ge-3/0/2 detail**

```
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
```



```
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 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
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Drop   bytes  :                   0                    0 bps
   Drop   packets:                   0                    0 pps
  Ingress 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

  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

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
```



```
   Output packets:                    0                    0 pps
  Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
    Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Local statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
  Transit statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
```

show interfaces extensive (Gigabit Ethernet IQ2)

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4004000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
  Last flapped   : 2007-11-07 21:31:41 PST (02:03:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :          38910844056                 7952 bps
   Output bytes  :              7174605                 8464 bps
   Input  packets:            418398473                   11 pps
   Output packets:                78903                   12 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  :          38910799145                 7952 bps
   Input  packets:            418397956                   11 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: 1, 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           418390823            418390823                   0

  1 expedited-fo                  0                    0                   0

  2 assured-forw                  0                    0                   0

  3 network-cont               7133                 7133                   0

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

  0 best-effort                1031                 1031                   0

  1 expedited-fo                  0                    0                   0

  2 assured-forw                  0                    0                   0

  3 network-cont              77872                77872                   0

Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                   38910844056          7174605
  Total packets                    418398473            78903
  Unicast packets               408021893366             1026
  Broadcast packets                       10               12
  Multicast packets                418398217            77865
  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 Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count               418398473
  Input packet rejects                   479
  Input DA rejects                       479
  Input SA rejects                         0
  Output packet count                                   78903
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
      Remote fault: OK
  Local resolution:
```



```
        Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 7
  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 ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             812400
     Output bytes  :            1349206
     Input  packets:               9429
     Output packets:               9449
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Local statistics:
     Input  bytes  :             812400
     Output bytes  :            1349206
     Input  packets:               9429
     Output packets:               9449
    Transit statistics:
     Input  bytes  :                  0              7440 bps
     Output bytes  :                  0              7888 bps
     Input  packets:                  0                10 pps
     Output packets:                  0                11 pps
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
      Output Filters: F2-ge-3/0/1.0-out (53)
      Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
        Generation: 196
    Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output 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 interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in Table 124 on page 1861.

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 6
    Protocol inet, MTU: 1500
      Flags: Unnumbered
      Donor interface: lo0.0 (Index 64)
      Preferred source address: 203.0.113.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
  Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-

    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
  Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
    Line Identity:
```



```
      Dynamic Profile: ali-set-profile
      Circuit-id Remote-id Accept-no-ids
    PPPoE:
      Dynamic Profile: ali-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      Direct Connect: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

Sample Output Gigabit Ethernet

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__
```

Sample Output

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
```

```
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
```

show interfaces detail (Gigabit Ethernet)

user@host> show interfaces ge-0/0/1 detail

```
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
 BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
```

```
 Interface transmit statistics: Disabled

 Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
   Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
   Traffic statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
   Local statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
   Transit statistics:
    Input  bytes  :                  0                    0 bps
    Output bytes  :                  0                    0 bps
    Input  packets:                  0                    0 pps
    Output packets:                  0                    0 pps
   Security: Zone: public
   Flow Statistics :
   Flow Input statistics :
     Self packets :                   0
     ICMP packets :                   0
     VPN packets :                    0
     Multicast packets :              0
     Bytes permitted by policy :      0
     Connections established :        0
   Flow Output statistics:
     Multicast packets :              0
     Bytes permitted by policy :      0
   Flow error statistics (Packets dropped due to):
     Address spoofing:                0
     Authentication failed:           0
     Incoming NAT errors:             0
     Invalid zone received packet:    0
     Multiple user authentications:   0
     Multiple incoming NAT:           0
     No parent for a gate:            0
     No one interested in self packets: 0
     No minor session:                0
     No more sessions:                0
     No NAT gate:                     0
     No route present:                0
     No SA for incoming SPI:          0
     No tunnel found:                 0
     No session for a gate:           0
     No zone or NULL zone binding     0
     Policy denied:                   0
     Security association not active:   0
     TCP sequence number out of window: 0
     Syn-attack protection:           0
     User authentication errors:      0
   Protocol inet, MTU: 1500, Generation: 150, Route table: 0
     Flags: Sendbcast-pkt-to-re
     Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
       Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150
```

show interfaces statistics st0.0 detail

```
user@host> show interfaces statistics st0.0 detail
  Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
    Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
    Traffic statistics:
     Input  bytes  :        528152756774
     Output bytes  :        575950643520
     Input  packets:         11481581669
     Output packets:         12520666095
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0          121859888 bps
     Output bytes  :                   0          128104112 bps
     Input  packets:                   0             331141 pps
     Output packets:                   0             348108 pps
    Security: Zone: untrust
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
    sap vrrp
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      0
      Multicast packets :                0
      Bytes permitted by policy :        525984295844
      Connections established :          7
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        576003290222
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  2000280
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 9192
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
 NH drop cnt: 0
    Generation: 155, Route table: 0
```



```
      Flags: Sendbcast-pkt-to-re
```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  MAC statistics:                      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: 2, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Incomplete
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0      low
none
    3 network-control         5       50000000     5              0      low
none
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Security: Zone: public
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      0
      Multicast packets :                0
      Bytes permitted by policy :        0
      Connections established :          0
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        0
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
```



```
    Authentication failed:             0
    Incoming NAT errors:               0
    Invalid zone received packet:      0
    Multiple user authentications:     0
    Multiple incoming NAT:             0
    No parent for a gate:              0
    No one interested in self packets: 0
    No minor session:                  0
    No more sessions:                  0
    No NAT gate:                       0
    No route present:                  0
    No SA for incoming SPI:            0
    No tunnel found:                   0
    No session for a gate:             0
    No zone or NULL zone binding       0
    Policy denied:                     0
    Security association not active:   0
    TCP sequence number out of window: 0
    Syn-attack protection:             0
    User authentication errors:        0
  Protocol inet, MTU: 1500, Generation: 150, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
      Generation: 150
```

show interfaces terse

```
user@host> show interfaces terse
    Interface               Admin Link Proto    Local                 Remote
    ge-0/0/0                up    up
    ge-0/0/0.0              up    up   inet     10.209.4.61/18
    gr-0/0/0                up    up
    ip-0/0/0                up    up
    st0                     up    up
    st0.1                   up    ready inet
    ls-0/0/0                up    up
    lt-0/0/0                up    up
    mt-0/0/0                up    up
    pd-0/0/0                up    up
    pe-0/0/0                up    up
    e3-1/0/0                up    up
    t3-2/0/0                up    up
    e1-3/0/0                up    up
    se-4/0/0                up    down
    t1-5/0/0                up    up
    br-6/0/0                up    up
    dc-6/0/0                up    up
    dc-6/0/0.32767          up    up
    bc-6/0/0:1              down  up
    bc-6/0/0:1.0            up    down
    dl0                     up    up
    dl0.0                   up    up   inet
    dsc                     up    up
    gre                     up    up
    ipip                    up    up
    lo0                     up    up
    lo0.16385               up    up   inet     10.0.0.1            --> 0/0
                                                10.0.0.16           --> 0/0
```



```
    lsi                     up    up
    mtun                    up    up
    pimd                    up    up
    pime                    up    up
    pp0                     up    up
```

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                          Admin Link
ce1-1/2/6                                                           up    up
e1-1/2/6                                                            up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                          Admin Link
ce1-1/2/3                                                           up    up

ds-1/2/3:1                                                          up    up

ds-1/2/3:2                                                          up    up
```

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                             Packets                  Bytes
      Destination class             (packet-per-second)    (bits-per-second)
                         gold                        0                      0
                                  (                  0) (                   0)
                       silver                        0                      0
                                  (                  0) (                   0)
  Logical interface so-0/1/3.0
                                             Packets                  Bytes
      Destination class             (packet-per-second)    (bits-per-second)
                         gold                        0                      0
                                  (                  0) (                   0)
```

```
                        silver                         0                    0
                              (                        0) (                 0)
```

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
```



```
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           iso
....
```

show interfaces flow-statistics (Gigabit Ethernet)

```
user@host> show interfaces flow-statistics ge-0/0/1.0
```

```
  Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 5161
    Output packets: 83
    Security: Zone: zone2
    Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf
pgm
    pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
 https ike
   netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
 xnm-ssl
    lsping
    Flow Statistics :
    Flow Input statistics :
      Self packets :                     0
      ICMP packets :                     0
      VPN packets :                      2564
      Bytes permitted by policy :        3478
      Connections established :          1
    Flow Output statistics:
      Multicast packets :                0
      Bytes permitted by policy :        16994
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:  0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination:  203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255
```

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
```

```
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH) (SRX devices)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
  19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:02-19:17:
.....
```

show interfaces load-balancing (SRX devices)

```
user@host> show interfaces load-balancing
Interface  State            Last change  Member count
ams0       Up               1d 00:50     2
ams1       Up               00:00:59     2
```

show interfaces load-balancing detail (SRX devices)

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface        : ams0
  State          : Up
  Last change    : 1d 00:51
  Member count   : 2
  Members        :
      Interface    Weight   State
      mams-2/0/0    10        Active
      mams-2/1/0    10        Active
```

show interfaces mac-database (All MAC Addresses on a Port SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames     Input bytes  Output frames    Output bytes
  00:00:00:00:00:00             1              56              0               0
  00:00:c0:01:01:02       7023810       323095260              0               0
  00:00:c0:01:01:03       7023810       323095260              0               0
  00:00:c0:01:01:04       7023810       323095260              0               0
  00:00:c0:01:01:05       7023810       323095260              0               0
  00:00:c0:01:01:06       7023810       323095260              0               0
  00:00:c0:01:01:07       7023810       323095260              0               0
  00:00:c0:01:01:08       7023809       323095214              0               0
  00:00:c0:01:01:09       7023809       323095214              0               0
  00:00:c0:01:01:0a       7023809       323095214              0               0
  00:00:c0:01:01:0b       7023809       323095214              0               0
  00:00:c8:01:01:02      30424784      1399540064       37448598      1722635508
  00:00:c8:01:01:03      30424784      1399540064       37448598      1722635508
  00:00:c8:01:01:04      30424716      1399536936       37448523      1722632058
  00:00:c8:01:01:05      30424789      1399540294       37448598      1722635508
  00:00:c8:01:01:06      30424788      1399540248       37448597      1722635462
  00:00:c8:01:01:07      30424783      1399540018       37448597      1722635462
  00:00:c8:01:01:08      30424783      1399540018       37448596      1722635416
  00:00:c8:01:01:09       8836796       406492616        8836795       406492570
  00:00:c8:01:01:0a      30424712      1399536752       37448521      1722631966
  00:00:c8:01:01:0b      30424715      1399536890       37448523      1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on a Service SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames     Input bytes  Output frames    Output bytes
  00:00:00:00:00:00             1              56              0               0
```

```
  00:00:c0:01:01:02        7023810      323095260            0              0
  00:00:c0:01:01:03        7023810      323095260            0              0
  00:00:c0:01:01:04        7023810      323095260            0              0
  00:00:c0:01:01:05        7023810      323095260            0              0
  00:00:c0:01:01:06        7023810      323095260            0              0
  00:00:c0:01:01:07        7023810      323095260            0              0
  00:00:c0:01:01:08        7023809      323095214            0              0
  00:00:c0:01:01:09        7023809      323095214            0              0
  00:00:c0:01:01:0a        7023809      323095214            0              0
  00:00:c0:01:01:0b        7023809      323095214            0              0
  00:00:c8:01:01:02       31016568     1426762128     38040381     1749857526
  00:00:c8:01:01:03       31016568     1426762128     38040382     1749857572
  00:00:c8:01:01:04       31016499     1426758954     38040306     1749854076
  00:00:c8:01:01:05       31016573     1426762358     38040381     1749857526
  00:00:c8:01:01:06       31016573     1426762358     38040381     1749857526
  00:00:c8:01:01:07       31016567     1426762082     38040380     1749857480
  00:00:c8:01:01:08       31016567     1426762082     38040379     1749857434
  00:00:c8:01:01:09        9428580      433714680      9428580      433714680
  00:00:c8:01:01:0a       31016496     1426758816     38040304     1749853984
  00:00:c8:01:01:0b       31016498     1426758908     38040307     1749854122
```

show interfaces mac-database mac-address

user@host> **show interfaces mac-database xe-0/3/3 mac-address (SRX devices) 00:00:c8:01:01:09**

```
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes      : 202324652
    Output bytes     : 202324560
    Input frames     : 4398362
    Output frames    : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

show interfaces mc-ae (SRX devices)

user@host> **show interfaces mc-ae ae0 unit 512**

```
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
     BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

show interfaces policers (SRX devices)

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
...
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
...
```

show interfaces policers interface-name (SRX devices)

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

show interfaces queue (SRX devices)

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```
user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0
  Queued:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
  Transmitted:
    Packets              :                     0                     0 pps
    Bytes                :                     0                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
     Low                 :                     0                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
     Low                 :                     0                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
  Queue Buffer Usage:
    Reserved buffer      :             118750000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 1, Forwarding classes: class1
..
..
 Queue Buffer Usage:
    Reserved buffer      :                  9192 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
    Reserved buffer      :               6250000 bytes
  Queue-depth bytes      :
    Current              :                     0
..
..
```

show interfaces redundancy (SRX devices)

```
user@host> show interfaces redundancy
Interface   State          Last change  Primary    Secondary  Current status
rsp0        Not present                 sp-1/0/0   sp-0/2/0   both down
rsp1        On secondary   1d 23:56     sp-1/2/0   sp-0/3/0   primary down
rsp2        On primary     10:10:27     sp-1/3/0   sp-0/2/0   secondary down
rlsq0       On primary     00:06:24     lsq-0/3/0  lsq-1/0/0  both up
```


show interfaces redundancy (Aggregated Ethernet SRX devices)

```
user@host> show interfaces redundancy
Interface  State          Last change  Primary     Secondary    Current status
rlsq0      On secondary    00:56:12     lsq-4/0/0    lsq-3/0/0    both up

ae0
ae1
ae2
ae3
ae4
```

show interfaces redundancy detail (SRX devices)

```
user@host> show interfaces redundancy detail
Interface        : rlsq0
  State          : On primary
  Last change    : 00:45:47
  Primary        : lsq-0/2/0
  Secondary      : lsq-1/2/0
  Current status : both up
  Mode           : hot-standby

Interface        : rlsq0:0
  State          : On primary
  Last change    : 00:45:46
  Primary        : lsq-0/2/0:0
  Secondary      : lsq-1/2/0:0
  Current status : both up
  Mode           : warm-standby
```

show interfaces routing brief (SRX devices)

```
user@host> show interfaces routing brief
Interface        State Addresses
so-5/0/3.0       Down  ISO   enabled
so-5/0/2.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.120
                       INET  enabled
so-5/0/1.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.130
                       INET  enabled
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

show interfaces routing detail (SRX devices)

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
```



```
  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...
```

show interfaces routing-instance all (SRX devices)

```
user@host> show interfaces terse routing-instance all
Interface    Admin    Link    Proto     Local              Remote Instance
at-0/0/1     up       up      inet      10.0.0.1/24
ge-0/0/0.0   up       up      inet       192.168.4.28/24          sample-a
at-0/1/0.0   up       up      inet6      fe80::a:0:0:4/64         sample-b
so-0/0/0.0   up       up      inet      10.0.0.1/32
```

show interfaces snmp-index (SRX devices)

```
user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P
```

show interfaces source-class all (SRX devices)

```
user@host> show interfaces source-class all
Logical interface so-0/1/0.0
                                              Packets                Bytes
      Source class                 (packet-per-second)    (bits-per-second)
                             gold              1928095            161959980
```

```
                                (             889) (             597762)
                         bronze                 0                      0
                                (               0) (                   0)
                         silver                 0                      0
                                (               0) (                   0)
  Logical interface so-0/1/3.0
                                          Packets                  Bytes
      Source class               (packet-per-second)    (bits-per-second)
                           gold                 0                      0
                                (               0) (                   0)
                         bronze                 0                      0
                                (               0) (                   0)
                         silver            116113                9753492
                                (             939) (             631616)
```

show interfaces statistics (Fast Ethernet SRX devices)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary, DCU, SCU-in
                                          Packets                  Bytes
      Destination class          (packet-per-second)    (bits-per-second)
                        silver1                 0                      0
                                (               0) (                   0)
                        silver2                 0                      0
                                (               0) (                   0)
                        silver3                 0                      0
                                (               0) (                   0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

show interfaces switch-port (SRX devices)

```
user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
        Speed: 100mbps, Auto-negotiation: Enabled
   Statistics:                         Receive         Transmit
          Total bytes                  28437086         21792250
          Total packets                  409145            88008
```



```
          Unicast packets                      9987            83817
          Multicast packets                  145002                0
          Broadcast packets                  254156             4191
          Multiple collisions                    23               10
          FIFO/CRC/Align errors                   0                0
          MAC pause frames                        0                0
          Oversized frames                        0
          Runt frames                             0
          Jabber frames                           0
          Fragment frames                         0
          Discarded frames                        0
        Autonegotiation information:
          Negotiation status: Complete
          Link partner:
              Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
 partner Speed: 100 Mbps
          Local resolution:
              Flow control: None, Remote fault: Link OK
```

show interfaces transport pm (SRX devices)

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                 COUNT            THRESHOLD        TCA-ENABLED      TCA-RAISED

  OTU-BBE                 0                800               No               No
  OTU-ES                  0                135               No               No
  OTU-SES                 0                90                No               No
  OTU-UAS                 427              90                No               No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                 COUNT            THRESHOLD        TCA-ENABLED      TCA-RAISED

  OTU-BBE                 0                800               No               No
  OTU-ES                  0                135               No               No
  OTU-SES                 0                90                No               No
  OTU-UAS                 0                90                No               No
  Near End       Suspect Flag:False             Reason:None
    PM                 COUNT            THRESHOLD        TCA-ENABLED      TCA-RAISED

  ODU-BBE                 0                800               No               No
  ODU-ES                  0                135               No               No
  ODU-SES                 0                90                No               No
  ODU-UAS                 427              90                No               No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                 COUNT            THRESHOLD        TCA-ENABLED      TCA-RAISED

  ODU-BBE                 0                800               No               No
  ODU-ES                  0                135               No               No
  ODU-SES                 0                90                No               No
  ODU-UAS                 0                90                No               No
  FEC            Suspect Flag:False             Reason:None
    PM                 COUNT            THRESHOLD        TCA-ENABLED      TCA-RAISED

  FEC-CorrectedErr        2008544300       0                 NA               NA
  FEC-UncorrectedWords    0                0                 NA               NA
  BER            Suspect Flag:False             Reason:None
    PM               MIN        MAX        AVG        THRESHOLD        TCA-ENABLED
     TCA-RAISED
  BER                3.6e-5     5.8e-5     3.6e-5      10.0e-3          No
```

```
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                           CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                           (MIN)
(MAX)    (MIN)  (MAX)        (MIN)  (MAX)
  Lane chromatic dispersion        0        0         0         0         0
 0        NA     NA          NA     NA
  Lane differential group delay    0        0         0         0         0
 0        NA     NA          NA     NA
    q Value                          120      120       120       120       0
   0        NA     NA          NA     NA
  SNR                               28       28        29        28        0
 0        NA     NA          NA     NA
  Tx output power(0.01dBm)         -5000    -5000     -5000     -5000     -300
 -100     No     No          No     No
  Rx input power(0.01dBm)          -3642    -3665     -3626     -3637     -1800
 -500     No     No          No     No
  Module temperature(Celsius)       46       46        46        46        -5
 75       No     No          No     No
  Tx laser bias current(0.1mA)     0        0         0         0         0
 0        NA     NA          NA     NA
  Rx laser bias current(0.1mA)     1270     1270      1270      1270      0
 0        NA     NA          NA     NA
  Carrier frequency offset(MHz)    -186     -186      -186      -186      -5000
 5000     No     No          No     No
```

show security zones (SRX devices)

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```


show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)

List of Syntax

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface) on page 1935
Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface) on page 1935

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface)

show interfaces em0 | fxp0
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface)

show interfaces bcm0 | em0 | em1| fxp1 | fxp2 | ixgbe0 | ixgbe1
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Description

(M Series, T Series, TX Matrix Plus, and PTX Series devices only) Display status information about the management Ethernet and internal Ethernet interfaces.

Options

em0 | fxp0—(M Series, MX Series, T Series, and PTX Series) Display standard information about the management Ethernet interface. For supported Ethernet interface by chassis and Routing Engine, see *Supported Routing Engines by Router*.

bcm0 | em0 | em1 | fxp1 | fxp2 | ixgbe0 | ixgbe1—(M Series, MX Series, T Series, and PTX Series) Display standard information about the internal Ethernet interfaces. See *Supported Routing Engines by Router* for the internal Ethernet interface names for each Routing Engine by hardware platform.

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

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information.

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

Output Fields Table 127 on page 1936 lists the output fields for the **show interfaces** (management) command on the M Series routers, T Series routers, TX Matrix Plus routers, and PTX Series. Output fields are listed in the approximate order in which they appear.

Table 127: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface

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	Physical interface index number, 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**
Type	Type of interface.	All levels

Table 127: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (continued)

Field Name	Field Description	Level of Output
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU)—Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	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 type	Data transmission type.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	**detail extensive**
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive**
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 packets	Number of packets received on the physical interface.	None specified
Output packets	Number of packets transmitted on the physical interface.	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 logical and physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**

Table 127: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (continued)

Field Name	Field Description	Level of Output
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	• **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, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of 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 dropped by the ASIC RED mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the “Device Flags” section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	**detail extensive** none
inet	IP address of the logical interface.	**brief**
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 127: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (continued)

Field Name	Field Description	Level of Output
Route table	Route table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about address flags. 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.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp0 brief
Physical interface: fxp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.70.143/21
```

show interfaces (Management Ethernet)

```
user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Last flapped   : Never
    Input packets : 80804
    Output packets: 1105

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255
```

show interfaces (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Last flapped   : Never
    Input packets : 1424
    Output packets: 5282

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 1424
    Output packets: 5282
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127
```

show interfaces (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Last flapped   : Never
 Input packets : 212581
 Output packets: 71

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 212551
 Output packets: 71
 Protocol inet, MTU: 1500
   Flags: Is-Primary
   Addresses, Flags: Is-Default Is-Preferred Is-Primary
     Destination: 192.168.3/24, Local: 192.168.3.30,
     Broadcast: 192.168.3.255
```

show interfaces detail (Management Ethernet)

```
user@host> show interfaces fxp0 detail
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
```



```
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6484031
   Output bytes  :               167503
   Input  packets:                81008
   Output packets:                 1110

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces detail (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 detail
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               124351
   Output bytes  :              1353212
   Input  packets:                 1804
   Output packets:                 5344
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Local statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
```

```
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1
```

show interfaces detail (Management Ethernet [PTX Packet Transport Routers])

```
user@host> show interfaces detail em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
     Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :           15255909
       Output bytes  :               4608
       Input  packets:             214753
       Output packets:                 72
       IPv6 transit statistics:
 Input  bytes  :                 0
 Output bytes  :                 0
 Input  packets:                 0
 Output packets:                 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
  Input  bytes  :           14394630
  Output bytes  :               3024
  Input  packets:             214723
  Output packets:                 72
 Local statistics:
  Input  bytes  :           14394630
  Output bytes  :               3024
  Input  packets:             214723
  Output packets:                 72
 Protocol inet, MTU: 1500, Generation: 1, Route table: 0
   Flags: Is-Primary
   Addresses, Flags: Is-Default Is-Preferred Is-Primary
     Destination: 192.168.3/24, Local: 192.168.3.30,
     Broadcast: 192.168.3.255, Generation: 1
```

show interfaces extensive (Management Ethernet)

```
user@host> show interfaces fxp0 extensive
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
```



```
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             6678904
   Output bytes  :              169657
   Input  packets:               83946
   Output packets:                1127
  Input errors:
    Errors: 12, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces extensive (Management Ethernet [TX Matrix Plus Router])

user@host> **show interfaces em0 extensive**

```
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              127120
   Output bytes  :             1357414
   Input  packets:                1843
   Output packets:                5372
   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, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0
```



```
  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :              119748
     Output bytes  :             1335719
     Input  packets:                1843
     Output packets:                5372
    Local statistics:
     Input  bytes  :              119748
     Output bytes  :             1335719
     Input  packets:                1843
     Output packets:                5372
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1
```

show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces extensive em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
     Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :            15236459
       Output bytes  :                4608
       Input  packets:              214482
       Output packets:                  72
       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, Giants: 0,
 Policed discards: 0, Resource errors: 0
      Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
  Input  bytes  :            14376264
  Output bytes  :                3024
  Input  packets:              214452
  Output packets:                  72
 Local statistics:
  Input  bytes  :            14376264
```



```
  Output bytes  :                 3024
  Input  packets:               214452
  Output packets:                   72
 Protocol inet, MTU: 1500, Generation: 1, Route table: 0
   Flags: Is-Primary
   Addresses, Flags: Is-Default Is-Preferred Is-Primary
     Destination: 192.168.3/24, Local: 192.168.3.30,
     Broadcast: 192.168.3.255, Generation: 1
```

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp1 brief
Physical interface: fxp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp1.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.0.0.4/8
    inet6 fe80::200:ff:fe00:4/64
          fec0::10:0:0:4/64
    tnp   4
```

show interfaces brief (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface em0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.178.11/25
```

show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
     Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps

      Logical interface em0.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet  192.168.3.30/24

    root@aboslutely> show interfaces em0 terse
    Interface               Admin Link Proto    Local                 Remote
    em0                     up    up
    em0.0                   up    up   inet     192.168.3.30/24
```

show interfaces (Internal Ethernet)

```
user@host> show interfaces fxp1
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 30655
    Output packets: 33323

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
      Addresses
        Local: 4
```

show interfaces (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 2301738
    Output packets: 3951155

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 2301595
    Output packets: 3951155
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 192.168.0.4
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
```

```
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
      Addresses
        Local: 0x22000004
```

show interfaces detail (Internal Ethernet)

```
user@host> show interfaces fxp1 detail
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2339969
   Output bytes  :             15880707
   Input  packets:                30758
   Output packets:                33443

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces detail (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 detail
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
```

```
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :           238172825
   Output bytes  :          1338948955
   Input  packets:             2360984
   Output packets:             4061512
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :           228720309
     Output bytes  :          1261387447
     Input  packets:             2360841
     Output packets:             4061512
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :           228720309
     Output bytes  :          1261387447
     Input  packets:             2360841
     Output packets:             4061512
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
       Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation:
 2
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255,
 Generation: 3
    Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
    Generation: 4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500, Generation: 5
    Generation: 4, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6
```

show interfaces extensive (internal Ethernet)

```
user@host> show interfaces fxp1 extensive
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
```



```
Speed: 100mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
Link type      : Full-Duplex
Physical info  : Unspecified
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :              2349897
 Output bytes  :             15888605
 Input  packets:                30896
 Output packets:                33607
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500, Generation: 7, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
      Generation: 3
  Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
      Broadcast: Unspecified, Generation: 5
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
      Generation: 7
  Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
    Flags: Primary, Is-Primary
    Addresses, Flags: None
      Destination: Unspecified, Local: 4, Broadcast: Unspecified,
      Generation: 8
```

show interfaces extensive (internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 extensive
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            242730780
```

```
   Output bytes  :          1348312269
   Input  packets:             2398737
   Output packets:             4133510
   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, Giants: 0, Policed discards:
0, Resource errors: 0
 Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

 Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
   Flags: SNMP-Traps Encapsulation: ENET2
   Traffic statistics:
    Input  bytes  :           233127252
    Output bytes  :          1269350897
    Input  packets:             2398594
    Output packets:             4133510
    IPv6 transit statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
   Local statistics:
    Input  bytes  :           233127252
    Output bytes  :          1269350897
    Input  packets:             2398594
    Output packets:             4133510
   Protocol inet, MTU: 1500, Generation: 2, Route table: 1
     Flags: Is-Primary
     Addresses, Flags: Is-Preferred
      Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation:
2
     Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
       Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255,
Generation: 3
   Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
     Flags: Is-Primary
     Addresses, Flags: Is-Preferred
       Destination: fe80::/64, Local: fe80::200:ff:fe22:4
   Generation: 4
     Addresses, Flags: Is-Default Is-Preferred Is-Primary
       Destination: fec0::/64, Local: fec0::a:22:0:4
   Protocol tnp, MTU: 1500, Generation: 5
   Generation: 4, Route table: 1
     Flags: Primary, Is-Primary
     Addresses, Flags: None
       Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6
```

show interfaces (PPPoE)

Syntax

show interfaces pp0.*logical*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

Options

pp0.*logical*—Display standard status information about the PPPoE 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 PPPoE interfaces.

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

statistics—(Optional) Display PPPoE interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (PPPoE) on page 1957
show interfaces (PPPoE over Aggregated Ethernet) on page 1957
show interfaces brief (PPPoE) on page 1958
show interfaces detail (PPPoE) on page 1958
show interfaces extensive (PPPoE on M120 and M320 Routers) on page 1959

Output Fields

Table 128 on page 1951 lists the output fields for the **show interfaces** (PPPoE) command. Output fields are listed in the approximate order in which they appear.

Table 128: show interfaces (PPPoE) 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	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 128: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
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**
Type	Physical interface type (PPPoE).	All levels
Link-level type	Encapsulation on the physical interface (PPPoE).	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be **Internal** or **External**.	All levels
Speed	Speed at which the interface is running.	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 type	Physical interface link type: **full duplex** or **half duplex**.	All levels
Link flags	Information about the interface. Possible values are described in the “Link Flags” section under *Common Output Fields Description*.	All levels
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
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive**
Hardware address	MAC address of the hardware.	**detail extensive**
Alternate link address	Backup address of the link.	**detail extensive**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 128: show interfaces (PPPoE) 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. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. NOTE: These fields include dropped traffic and exception traffic, as those fields are not separately defined. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of 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. • **Giants**—Number of frames received that are larger than the giant 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. • **Resource errors**—Sum of B chip Tx drops and IXP Tx net transmit drops.	**extensive**
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), then the cable, the far-end system, or the 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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of B chip Tx drops and IXP Tx net transmit drops.	**extensive**
Logical Interface		

Table 128: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number (which reflects its initialization sequence).	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**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
Encapsulation	Type of encapsulation configured on the logical interface.	All levels
PPP parameters	PPP status: • LCP restart timer—Length of time (in milliseconds) between successive Link Control Protocol (LCP) configuration requests. • NCP restart timer—Length of time (in milliseconds) between successive Network Control Protocol (NCP) configuration requests.	**detail**
PPPoE	PPPoE status: • **State**—State of the logical interface (**up** or **down**). • **Session ID**—PPPoE session ID. • **Service name**—Type of service required. Can be used to indicate an Internet service provider (ISP) name or a class or quality of service. • **Configured AC name**—Configured access concentrator name. • **Auto-reconnect timeout**—Time after which to try to reconnect after a PPPoE session is terminated, in seconds. • **Idle Timeout**—Length of time (in seconds) that a connection can be idle before disconnecting. • **Underlying interface**—Interface on which PPPoE is running.	All levels
Link	Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	**detail extensive**

Table 128: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router.	**detail extensive**
Keepalive settings	(PPP and HDLC) Configured settings for keepalives. • **interval** ***seconds***—The time in seconds between successive keepalive requests. The range is **10** seconds through **32,767** seconds, with a default of **10** seconds. • **down-count** ***number***—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is **1** through **255**, with a default of **3**. • **up-count** ***number***—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is **1** through **255**, with a default of **1**.	**detail extensive**

Keepalive statistics	(PPP and HDLC) Information about keepalive packets.
	
	• **Input**—Number of keepalive packets received by PPP.
	  ◦ **(last seen 00:00:00 ago)**—Time the last keepalive packet was received, in the format ***hh:mm:ss***.
	• **Output**—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ◦ **(last seen 00:00:00 ago)**—Time the last keepalive packet was sent, in the format ***hh:mm:ss***.
	
	(MX Series routers with MPCs/MICs) When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the display does not include the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.
		detail extensive

Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

Table 128: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
LCP state	(PPP) Link Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—LCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—LCP is not configured on the interface. • **Opened**—LCP negotiation is successful.	none **detail extensive**
NCP state	(PPP) Network Control Protocol state. • **Conf-ack-received**—Acknowledgement was received. • **Conf-ack-sent**—Acknowledgement was sent. • **Conf-req-sent**—Request was sent. • **Down**—NCP negotiation is incomplete (not yet completed or has failed). • **Not-configured**—NCP is not configured on the interface. • **Opened**—NCP negotiation is successful.	**detail extensive** none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. • **Chap-Chal-received**—Challenge was received but response not yet sent. • **Chap-Chal-sent**—Challenge was sent. • **Chap-Resp-received**—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • **Chap-Resp-sent**—Response was sent for the challenge received. • **Closed**—CHAP authentication is incomplete. • **Failure**—CHAP authentication failed. • **Not-configured**—CHAP is not configured on the interface. • **Success**—CHAP authentication was successful.	none **detail extensive**
Protocol	Protocol family configured on the logical interface.	**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**
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing 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 the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none

Table 128: show interfaces (PPPoE) Output Fields (continued)

Field Name	Field Description	Level of Output
Addresses, Flags	Information about the addresses configured for the protocol family. 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.	**detail extensive** none

Sample Output

show interfaces (PPPoE)

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type       : Full-Duplex
  Link flags      : None
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)

  Logical interface pp0.0 (Index 72) (SNMP ifIndex 72)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
  Input packets : 0
  Output packets: 0
  LCP state: Not-configured
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
    Protocol inet, MTU: 100
      Flags: User-MTU, Negotiate-Address
```

show interfaces (PPPoE over Aggregated Ethernet)

```
user@host> show interfaces pp0.1073773821
Logical interface pp0.1073773821 (Index 80) (SNMP ifIndex 32584)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: alcor, Remote MAC address: 00:00:5e:00:53:01,
      Underlying interface: demux0.100 (Index 88)
    Link:
      ge-1/0/0.32767
      ge-1/0/1.32767
```



```
    Input packets : 6
    Output packets: 6
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,  mpls:
Not-configured
  CHAP state: Closed
  PAP state: Success
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Primary
        Local: 203.0.113.1
```

show interfaces brief (PPPoE)

```
user@host> show interfaces pp0 brief
Physical interface: pp0, Enabled, Physical link is Up
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface pp0.0
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
    inet
```

show interfaces detail (PPPoE)

```
user@host> show interfaces pp0 detail
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 24, Generation: 9
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  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
  Logical interface pp0.0 (Index 72) (SNMP ifIndex 72) (Generation 14)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Service name: None, Configured AC name: sapphire,
      Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
      Underlying interface: at-5/0/0.0 (Index 70)
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
```

```
      Output packets:                   0
     Local statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
     Transit statistics:
      Input  bytes  :                   0                    0 bps
      Output bytes  :                   0                    0 bps
      Input  packets:                   0                    0 pps
      Output packets:                   0                    0 pps
   LCP state: Not-configured
   NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
   mpls: Not-configured
   CHAP state: Closed
     Protocol inet, MTU: 100, Generation: 14, Route table: 0
       Flags: User-MTU, Negotiate-Address
```

show interfaces extensive (PPPoE on M120 and M320 Routers)

```
user@host> show interfaces pp0 extensive
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 93, Generation: 129
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              972192                    0 bps
   Output bytes  :              975010                    0 bps
   Input  packets:                1338                    0 pps
   Output packets:                1473                    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, Giants: 0, Policed discards:
 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0

  Logical interface pp0.0 (Index 69) (SNMP ifIndex 96) (Generation 194)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 26,
      Session AC name: None, AC MAC address: 00:00:5e:00:53:12,
      Service name: None, Configured AC name: None,
      Auto-reconnect timeout: Never, Idle timeout: Never,
      Underlying interface: ge-3/0/1.0 (Index 67)
    Traffic statistics:
     Input  bytes  :                 252
```

```
      Output bytes  :                 296
      Input  packets:                   7
      Output packets:                   8
      IPv6 transit statistics:
       Input  bytes  :                  0
       Output bytes  :                  0
       Input  packets:                  0
       Output packets:                  0
     Local statistics:
      Input  bytes  :                 252
      Output bytes  :                 296
      Input  packets:                   7
      Output packets:                   8
     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
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 1 (last seen 00:00:00 ago)
    Output: 1 (last sent 00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Closed
    Protocol inet, MTU: 1492, Generation: 171, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
Generation: 206
```


show interfaces interface-set (Ethernet Interface Set)

Syntax

show interfaces interface-set *interface-set-name*
<detail | terse>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Display information about the specified gigabit or 10-Gigabit Ethernet interface set.

You can also use the **show interfaces interface-set** command to display information about agent circuit identifier (ACI) interface sets.

Options

interface-set ***interface-set-name***—Display information about the specified Gigabit Ethernet, 10-Gigabit Ethernet, ACI, or ALI interface set.

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

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

List of Sample Output

show interfaces interface-set terse on page 1963
show interfaces interface-set detail on page 1963
show interfaces interface-set (ACI Interface Set based on ACI) on page 1963
show interfaces interface-set (ACI Interface Set based on ACI Trusted Option) on page 1964
show interfaces interface-set (ACI Interface Set based on ARI Trusted Option) on page 1964
show interfaces interface-set (ACI Interface Set based on ARI Trusted Option when both ACI and ARI are received) on page 1964
show interfaces interface-set (ACI Interface Set based on Accept-No-IDs Trusted Option when neither ACI nor ARI is received) on page 1964
show interfaces interface-set (L2BSA and PPPoE Subscribers) on page 1965

Output Fields

Table 129 on page 1961 describes the information for the **show interfaces interface-set** command. Output fields are listed in the approximate order in which they appear.

Table 129: Ethernet show interfaces interface-set Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 129: Ethernet show interfaces interface-set Output Fields (continued)

Field Name	Field Description	Level of Output
Interface set	Name of the interface set or sets. For ACI interface sets, the set name is prefixed with **aci-**. For ALI interface sets, the set name is prefixed with the trusted option that the interface set is based on: • **aci-**— The ACI is configured as the trusted option. • **ari-**— The ARI is configured as the trusted option. • **aci+ari-**— Both ACI and ARI are configured as the trusted option. • **noids-**— Neither the ACI nor the ARI is configured as the trusted option and neither ACI nor ARI is received.	All levels
Interface set index	Index number of the interface set.	**detail** none
ACI VLAN	For ACI interface sets, the string received in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node. Only the **Agent Circuit ID** can be used to create the interface set. NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the **line-identity** autoconfiguration stanza.	**detail** none
Line Identity	For ALI interface sets, the trusted option received in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node. The trusted option can be either or both of the following: • **Agent Circuit ID**—The ACI value • **Agent Remote ID**—The ARI value. NOTE: When only **accept-no-ids** is configured as the trusted option, this field is not displayed. NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the **agent-circuit-id** autoconfiguration stanza.	**detail** none
PPPoE	Dynamic PPPoE subscriber interface that the router creates using the ACI or ALI interface set.	**detail** none
Max Sessions	For dynamic PPPoE subscriber interfaces, maximum number of PPPoE logical interfaces that that can be activated on the underlying interface.	**detail** none
Max Sessions VSA Ignore	For dynamic PPPoE subscriber interfaces, whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configured with the **max-sessions** statement: **Off** (default) or **On**.	**detail** none

Table 129: Ethernet show interfaces interface-set Output Fields (continued)

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes and number of bytes per second received and transmitted on the interface set • **Input packets, Output packets**—Number of packets and number of packets per second received and transmitted on the interface set.	**detail**
Egress queues supported	Total number of egress queues supported on the specified interface set.	**detail**
Egress queues in use	Total number of egress queues used on the specified interface set.	**detail**
Queue counters	**Queued packets**, **Transmitted packets**, and **Dropped packets** statistics for the four forwarding classes.	**detail**
Members	List of all interface sets or, for ACI interface sets, list of all subscriber interfaces belonging to the specified ACI interface set.	**detail** none

Sample Output

show interfaces interface-set terse

```
user@host> show interfaces interface-set terse
Interface set:
   iflset-xe-11/3/0-0
   ge-1/0/1-0
   ge-1/0/1-2
```

show interfaces interface-set detail

```
user@host> show interfaces interface-set iflset-xe-11/3/0-0 detail
Interface set: iflset-xe-11/3/0-0
  Interface set index: 19
  Traffic statistics:
   Output bytes  :            751017840           401673504 bps
   Output packets:             11044380              738377 pps
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets     Dropped packets
    0                        211091327             11044380           199995746
    1                                0                    0                   0
    2                                0                    0                   0
    3                                0                    0                   0
  Members:
   xe-11/3/0.0
```

show interfaces interface-set (ACI Interface Set based on ACI)

```
user@host> show interfaces interface-set
Interface set: aci-1001-ge-5/2/0.10
  Interface set index: 1
```

```
  Interface set snmp index: 67108865
    ACI VLAN:
      Agent Circuit ID: circuit0
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225472
```

show interfaces interface-set (ACI Interface Set based on ACI Trusted Option)

```
user@host> show interfaces interface-set
Interface set: ari-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
    Line Identity:
      Agent Circuit ID: remote20
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225474
```

show interfaces interface-set (ACI Interface Set based on ARI Trusted Option)

```
user@host> show interfaces interface-set
Interface set: aci-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
    Line Identity:
      Agent Remote ID: remote20
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225474
```

show interfaces interface-set (ACI Interface Set based on ARI Trusted Option when both ACI and ARI are received)

```
user@host> show interfaces interface-set
Interface set: ari-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
    Line Identity:
      Agent Remote ID: remote20
    PPPoE:
      Max Sessions: 32000, Max Sessions VSA Ignore: Off
  Members:
   demux0.3221225474
```

show interfaces interface-set (ACI Interface Set based on Accept-No-IDs Trusted Option when neither ACI nor ARI is received)

```
user@host> show interfaces interface-set
Interface set: noids-1002-demux0.3221225473
  Interface set index: 2
  Interface set snmp index: 67108866
  Members:
   demux0.3221225474
```

show interfaces interface-set (L2BSA and PPPoE Subscribers)

```
user@host> show interfaces interface-set
Interface set: ge-1/0/4
  Interface set index: 6
  Members:
   ge-1/0/4.1073741908
   pp0.1073741907
```

show interfaces interface-set queue

Syntax

show interfaces interface-set queue *interface-set-name*
<aggregate | remaining-traffic>
<forwarding-class *class-name*>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Display information about the gigabit or 10-Gigabit Ethernet interface set queue. Supported in MX Series routers with enhanced queuing DPCs.

Options

interface-set-name—(Optional) Display information about the specified gigabit or 10-Gigabit Ethernet interface set. Wildcard values can be used in the interface set name.

aggregate—(Optional) Display the aggregated queuing statistics of all member logical interfaces for interface sets that have traffic-control profiles configured.

both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics.

egress—(Optional) Display egress queue statistics.

forwarding-class *class-name*—(Optional) Display queuing statistics for the specified forwarding class.

ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics.

remaining-traffic—(Optional) Display the queuing statistics of all member logical interfaces for interface sets that do not have traffic-control profiles configured.

Required Privilege Level

view

List of Sample Output

show interfaces interface-set queue (Gigabit Ethernet) on page 1967
show interfaces interface-set queue both-ingress-egress (Enhanced DPC) on page 1968
show interfaces interface-set queue egress (Enhanced DPC) on page 1970
show interfaces interface-set queue forwarding-class (Gigabit Ethernet) on page 1971
show interfaces interface-set queue (Enhanced DPC) on page 1972
show interfaces interface-set queue remaining-traffic (Gigabit Ethernet) on page 1973

Output Fields

Table 130 on page 1966 describes the information for the **show interfaces interface-set queue** command.

Table 130: Ethernet show interfaces interface-set queue Output Fields

Field Name	Field Description	Level of Output
Physical Interface		

Table 130: Ethernet show interfaces interface-set queue Output Fields (continued)

Field Name	Field Description	Level of Output
Interface set	Name of the interface set.	All levels
Interface set index	Index number of the interface set.	All levels
Forwarding classes supported	Total number of forwarding classes supported on the specified interface set.	All levels
Forwarding classes in use	Total number of forwarding classes used on the specified interface set.	All levels
Egress queues supported	Total number of egress queues supported on the specified interface set.	All levels
Egress queues in use	Total number of egress queues used on the specified interface set.	All levels
Ingress queues supported	Total number of ingress queues supported on the specified interface set.	All levels
Ingress queues in use	Total number of ingress queues used on the specified interface set.	All levels
Queue	Egress or ingress queue number for the statistics being displayed.	All levels
Forwarding classes	Forwarding class name for the statistics being displayed.	All levels
Queued	**Packet** and **Byte** statistics for the specified queue. • **Packets**—Number of packets queued and input rate in packets per second. • **Bytes**—Number of bytes queued and input rate in bytes per second.	All levels
Transmitted	**Packet** and **Byte** statistics for the specified forwarding class. • **Packets**—Number of packets transmitted and transmit rate in packets per second. • **Bytes**—Number of bytes transmitted and transmit rate in bytes per second. • **Tail-dropped packets**—Number of packets tail dropped. • **RED-dropped packets**—Number of RED-dropped packets for the **low**, **medium-low**, **medium-high**, and **high** loss priorities. • **RED-dropped bytes**—Number of RED-dropped bytes for the **low**, **medium-low**, **medium-high**, and **high** loss priorities.	All levels

Sample Output

show interfaces interface-set queue (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0
Interface set: ge-2/2/0-0
  Interface set index: 3
```

```
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              3998482                    1 pps
    Bytes                :            271896884                  688 bps
  Transmitted:
    Packets              :              1077474                    1 pps
    Bytes                :             73268340                  688 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :              2921008                    0 pps
     Low                 :              2921008                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :            198628544                    0 bps
     Low                 :            198628544                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
...
```

show interfaces interface-set queue both-ingress-egress (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 both-ingress-egress
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            185968478               473161 pps
    Bytes                :          10042313520            204441336 bps
  Transmitted:
    Packets              :              5441673                13780 pps
    Bytes                :            293850342              5952960 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :            180526772               459372 pps
    RED-dropped bytes    :           9748446282            198451512 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :            522021472               473602 pps
    Bytes                :          28190332480            204599944 bps
  Transmitted:
    Packets              :              5791772                 4055 pps
    Bytes                :            312755688              1751976 bps
```



```
    Tail-dropped packets :                    0                  0 pps
    RED-dropped packets  :            516227139             469546 pps
    RED-dropped bytes    :          27876265560          202843872 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                  0 pps
    Bytes                :                    0                  0 bps
  Transmitted:
    Packets              :                    0                  0 pps
    Bytes                :                    0                  0 bps
    Tail-dropped packets :                    0                  0 pps
    RED-dropped packets  :                    0                  0 pps
    RED-dropped bytes    :                    0                  0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              5417304              13797 pps
    Bytes                :            368429508            7506096 bps
  Transmitted:
    Packets              :              5014996              12769 pps
    Bytes                :            341019728            6946560 bps
    Tail-dropped packets :                    0                  0 pps
    RED-dropped packets  :               402189               1028 pps
     Low                 :               402189               1028 pps
     Medium-low          :                    0                  0 pps
     Medium-high         :                    0                  0 pps
     High                :                    0                  0 pps
    RED-dropped bytes    :             27348852             559536 bps
     Low                 :             27348852             559536 bps
     Medium-low          :                    0                  0 bps
     Medium-high         :                    0                  0 bps
     High                :                    0                  0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                  0 pps
    Bytes                :                    0                  0 bps
  Transmitted:
    Packets              :                    0                  0 pps
    Bytes                :                    0                  0 bps
    Tail-dropped packets :                    0                  0 pps
    RED-dropped packets  :                    0                  0 pps
     Low                 :                    0                  0 pps
     Medium-low          :                    0                  0 pps
     Medium-high         :                    0                  0 pps
     High                :                    0                  0 pps
    RED-dropped bytes    :                    0                  0 bps
     Low                 :                    0                  0 bps
     Medium-low          :                    0                  0 bps
     Medium-high         :                    0                  0 bps
     High                :                    0                  0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :              5770534               3963 pps
    Bytes                :            396943252            2156144 bps
  Transmitted:
    Packets              :              3945152               1457 pps
    Bytes                :            268270336             792608 bps
    Tail-dropped packets :                    0                  0 pps
    RED-dropped packets  :              1815141               2506 pps
     Low                 :              1815141               2506 pps
```

```
      Medium-low          :                   0                 0 pps
      Medium-high         :                   0                 0 pps
      High                :                   0                 0 pps
     RED-dropped bytes    :           123429524           1363536 bps
      Low                 :           123429524           1363536 bps
      Medium-low          :                   0                 0 bps
      Medium-high         :                   0                 0 bps
      High                :                   0                 0 bps
 Queue: 3, Forwarding classes: network-control
   Queued:
     Packets              :                   0                 0 pps
     Bytes                :                   0                 0 bps
   Transmitted:
     Packets              :                   0                 0 pps
     Bytes                :                   0                 0 bps
     Tail-dropped packets :                   0                 0 pps
     RED-dropped packets  :                   0                 0 pps
      Low                 :                   0                 0 pps
      Medium-low          :                   0                 0 pps
      Medium-high         :                   0                 0 pps
      High                :                   0                 0 pps
     RED-dropped bytes    :                   0                 0 bps
      Low                 :                   0                 0 bps
      Medium-low          :                   0                 0 bps
      Medium-high         :                   0                 0 bps
      High                :                   0                 0 bps
```

show interfaces interface-set queue egress (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 egress
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             3958253             13822 pps
    Bytes                :           269217592           7519712 bps
  Transmitted:
    Packets              :             3665035             12729 pps
    Bytes                :           249222380           6924848 bps
    Tail-dropped packets :                   0                 0 pps
    RED-dropped packets  :              293091              1093 pps
     Low                 :              293091              1093 pps
     Medium-low          :                   0                 0 pps
     Medium-high         :                   0                 0 pps
     High                :                   0                 0 pps
    RED-dropped bytes    :            19930188            594864 bps
     Low                 :            19930188            594864 bps
     Medium-low          :                   0                 0 bps
     Medium-high         :                   0                 0 bps
     High                :                   0                 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                   0                 0 pps
    Bytes                :                   0                 0 bps
  Transmitted:
    Packets              :                   0                 0 pps
    Bytes                :                   0                 0 bps
    Tail-dropped packets :                   0                 0 pps
    RED-dropped packets  :                   0                 0 pps
```



```
         Low                 :                    0                    0 pps
         Medium-low          :                    0                    0 pps
         Medium-high         :                    0                    0 pps
         High                :                    0                    0 pps
        RED-dropped bytes    :                    0                    0 bps
         Low                 :                    0                    0 bps
         Medium-low          :                    0                    0 bps
         Medium-high         :                    0                    0 bps
         High                :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :              5350989                 3904 pps
    Bytes                :            368412924              2124048 bps
  Transmitted:
    Packets              :              3790469                 1465 pps
    Bytes                :            257751892               796960 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :              1550282                 2439 pps
     Low                 :              1550282                 2439 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :            105419176              1327088 bps
     Low                 :            105419176              1327088 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

show interfaces interface-set queue forwarding-class (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 forwarding-class best-effort
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            101857694              1420083 pps
    Bytes                :           6927234456            772532320 bps
  Transmitted:
    Packets              :              3984693                55500 pps
    Bytes                :            270959592             30192512 bps
    Tail-dropped packets :                    0                    0 pps
```

```
    RED-dropped packets  :             97870952              1364583 pps
     Low                 :             97870952              1364583 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :           6655225776            742339808 bps
     Low                 :           6655225776            742339808 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :                    0                    0 bps
```

show interfaces interface-set queue (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 ingress
Interface set: foo
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            149036817               473711 pps
    Bytes                :           8048003934            204642936 bps
  Transmitted:
    Packets              :              4360749                13891 pps
    Bytes                :            235480446              6000912 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :            144676035               459820 pps
    RED-dropped bytes    :           7812506592            198642024 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :            485089207               473605 pps
    Bytes                :          26195987476            204597576 bps
  Transmitted:
    Packets              :              5480799                 3959 pps
    Bytes                :            295963146              1710504 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :            479605853               469646 pps
    RED-dropped bytes    :          25898716170            202887072 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
```

show interfaces interface-set queue remaining-traffic (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 remaining-traffic
Interface set: ge-2/2/0-0
  Interface set index: 12
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :               2201552                     0 pps
    Bytes                :             149705536                     0 bps
  Transmitted:
    Packets              :                609765                     0 pps
    Bytes                :              41464020                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1591787                     0 pps
     Low                 :               1591787                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :             108241516                     0 bps
     Low                 :             108241516                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
```


show interfaces interval

Syntax

show interfaces interval
<*interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display the channel service unit (CSU) interface alarm and error count in 15-minute intervals for the past 24 hours. If the system has been operational for less than 24 hours, the maximum number of intervals available is displayed.

Options

interface-name—(Optional) Name of a particular interface.

Required Privilege Level

view

Related Documentation

- clear interfaces interval on page 1457

List of Sample Output

show interfaces interval (Channelized OC12) on page 1976
show interfaces interval (E3) on page 1976
show interfaces interval (SONET/SDH) on page 1976

Output Fields

Table 131 on page 1974 lists the output fields for the **show interfaces interval** command. Output fields are listed in the approximate order in which they appear.

Table 131: show interfaces interval Output Fields

Field Name	Field Description
Physical interface	Name of the interface.
SNMP ifIndex	SNMP index number for the physical interface.
***hh:mm*-current**	Time of day (in hours and minutes) at the beginning of the latest counter interval. The value of the latest counter interval is always less than 15 minutes.
hh:mm-hh:mm	Time of day (in hours and minutes) at the beginning and end of each 15-minute interval.
alarm or **event:** ***n***	Count of alarms and events within each 15-minute interval. The specific alarm or event depends on the interface media type. For a description of the alarm or event listed, see the ***interface-type*** **media** field (for example, **T1 media**) under the **show interfaces** command for the particular interface type in which you are interested.
Interval Total	Sum of all the alarm and defect counters for the last 24-hour period.

Table 131: show interfaces interval Output Fields (continued)

Field Name	Field Description
Interval Total	Sum of all the alarm and defect counters for the last 24-hour period.
Current Day Interval Total	Sum of all the alarm and defect counters in the current day. NOTE: The **Current Day Interval** output field is reset after 24 hours.
Previous Day Interval Total	Sum of all the alarm and defect counters in the previous day.

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-2/2/0
Physical interface: so-2/2/0, SNMP ifIndex: 553
  02:53-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  02:38-02:53:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  02:23-02:38:
```

```
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  02:08-02:23:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  01:53-02:08:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  01:38-01:53:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  01:23-01:38:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  01:08-01:23:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  00:53-01:08:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  00:38-00:53:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
....
  Current Day Interval Total:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
  Previous Day Interval Total (Last updated at 02:23):
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
```

show interfaces irb

Syntax

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

Release Information

Command introduced in Junos OS Release 8.4.

Description

Display integrated routing and bridging interfaces information.

Options

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

descriptions—(Optional) Display interface description strings.

mac—Display hardware MAC address

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

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

statistics—(Optional) Display static interface statistics.

Additional Information

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another bridging domain that has a Layer 3 protocol configured.

Required Privilege Level

view

List of Sample Output

show interfaces irb extensive on page 1982
show interfaces irb snmp-index on page 1983

Output Fields

Table 132 on page 1978 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 132: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels

Table 132: show interfaces irb Output Fields (continued)

Field Name	Field Description	Level of Output
Proto	Protocol configured on the interface.	**terse**
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Physical interface type.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**detail extensive brief** none
MTU	MTU size on the physical interface.	**detail extensive brief** none
Clocking	Reference clock source: **Internal** or **External**. Always unspecified on IRB interfaces.	**detail extensive brief**
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	**detail extensive brief**
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	**detail extensive brief** none
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	**detail extensive brief** none
Link type	Physical interface link type: **full duplex** or **half duplex**.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	**detail extensive** none
Physical Info	Physical interface information.	All levels
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	MAC address of the hardware.	**detail extensive** none
Alternate link address	Backup address of the link.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 132: show interfaces irb 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. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
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. • **Giants**—Number of frames received that are larger than the giant 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. • **Resource errors**—Sum of transmit drops.	**detail extensive**
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 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 DPC 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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**

Logical Interface

Table 132: show interfaces irb Output Fields (continued)

Field Name	Field Description	Level of Output
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 of 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*.	**detail extensive**
Encapsulation	Encapsulation on the logical interface.	**detail extensive**
Bandwidth	Dummy value that is ignored by an IRB interface. IRB interfaces are pseudo interfaces and do not have physical bandwidth associated with them.	**detail extensive**
Routing Instance	Routing instance IRB is configured under.	**detail extensive**
Bridging Domain	Bridging domain IRB is participating in.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical 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.	**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. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router.	**detail extensive**
Protocol	Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive**
MTU	Maximum transmission unit size on the logical interface.	**detail extensive**
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none

Table 132: show interfaces irb Output Fields (continued)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive**
Policer	The policer that is to be evaluated when packets are received or transmitted on the interface.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	**detail extensive**

Sample Output

show interfaces irb extensive

```
user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   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, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
```

```
    Routing Instance: customer_0 Bridging Domain: bd0
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     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  :                    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: 154, Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
        Generation: 155
    Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer
```

show interfaces irb snmp-index

```
user@host> show interfaces snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
    Protocol multiservice, MTU: 1500
      Flags: Is-Primary
```

show interfaces mac-database

Syntax

show interfaces mac-database (ge-*fpc/pic/port* | ge-*fpc/pic/port.n* | aex| xe-*fpc/pic/port* | xe-*fpc/pic/port.n* | et-*fpc/pic/port* | et-*fpc/pic/port.n*) <mac-address *mac-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.
Support for statement with the **ae*x*** option introduced in Junos OS Release 15.1.

Description

(M Series, T Series, MX Series routers, and PTX Series Packet Transport Routers only) Display media access control (MAC) address information for the specified interface.

Options

ge-*fpc/pic/port*—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

ge-*fpc/pic/port.n*—Display MAC addresses that have been learned on a particular logical interface.

ae*x*—Display MAC addresses that have been learned on a particular aggregated Ethernet interface.

xe-*fpc/pic/port*—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

xe-*fpc/pic/port.n*—Display MAC addresses that have been learned on a particular logical interface.

et-*fpc/pic/port*—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

et-*fpc/pic/port.n*—Display MAC addresses that have been learned on a particular logical interface.

mac-address *mac-address*—(Optional) Display detailed MAC address statistics, including policer information for **ge**, **xe**, and **et** interfaces.

Additional Information

On IQ2 PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

Required Privilege Level

view

List of Sample Output

show interfaces mac-database (All MAC Addresses on a Port) on page 1987
show interfaces mac-database (All MAC Addresses on an Aggregated Ethernet Interface) on page 1988
show interfaces mac-database (All MAC Addresses on a Service) on page 1989

show interfaces mac-database mac-address on page 1989

Output Fields Table 133 on page 1986 lists the output fields for the **show interfaces mac-database** command. Output fields are listed in the approximate order in which they appear.

Table 133: show interfaces mac-database Output Fields

Field Name	Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.
Interface index	Physical interface index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Description	Description and name of the interface.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback: **local** or **remote**.
Source filtering	Whether source filtering is configured.
Flow control	Whether flow control is enabled or disabled.
Minimum links needed	(Aggregated Ethernet interfaces only) Number of child links that must be operational for the aggregated interface to be operational.
Minimum bandwidth needed	(Aggregated Ethernet interfaces only) Minimum amount of bandwidth of child links that must be operational for the aggregated interface to be operational.
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.
Currrent address	(Aggregated Ethernet interfaces only) Configured MAC address.
Hardware address	(Aggregated Ethernet interfaces only) Hardware MAC address.
Last flapped	(Aggregated Ethernet interfaces only) Date, time, and how long ago the interface went from down to up or from up to down. The format is **Last flapped:** ***year-month-day hours:minutes:seconds timezone*** **(*wweeksddays hours:minutes* ago)**. For example, **Last flapped: 2013-12-18 04:33:22 PST (1w5d 22:23 ago).**

Table 133: show interfaces mac-database Output Fields (continued)

Field Name	Field Description
Input Rate	(Aggregated Ethernet interfaces only) Input rate in bits per second (bps) and packets per second (pps).
Output Rate	(Aggregated Ethernet interfaces only) Output rate in bps and pps.
Interface flags	Information about the interface. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.
Link flags	Information about the link. Possible v4alues are described in the “Device Flags” section under *Common Output Fields Description*.
Logical Interface	
Logical interface	Name of the logical interface.
Index	Logical interface index number, which reflects its initialization sequence.
SNMP ifIndex	Logical interface SNMP interface index number.
Flags	Information about the logical interface (possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.
Encapsulation	Encapsulation on the logical interface.
MAC address, Input frames, Input bytes, Output frames, Output bytes	MAC address and corresponding number of input frames, input bytes, output frames, and output bytes.
Number of MAC addresses	Number of MAC addresses configured.

Policer Statistics	(Displayed for **mac-address** option for **ge**, **xe**, and **et** interface types only) Display information about policers applied to a logical interface-MAC pair.
	• **Policer type**—Type of policer that is out of spec with respect to the configuration. It can be one or more of the following:
	  ◦ **Input premium**—Number of high-priority rating out-of-spec frames or bytes received.
	  ◦ **Output premium**—Number of high-priority rating out-of-spec frames or bytes sent.
	  ◦ **Input aggregate**—Total number of out-of-spec frames or bytes received.
	  ◦ **Output aggregate**—Total number of out-of-spec frames or bytes sent.
	• **Discarded Frames**—Number of discarded frames.
	• **Discarded Bytes**—Number of discarded bytes.

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

`user@host>` **show interfaces mac-database xe-0/3/3**

```
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address         Input frames    Input bytes  Output frames   Output bytes
  00:00:00:00:00:00              1             56              0              0
  00:00:c0:01:01:02        7023810      323095260              0              0
  00:00:c0:01:01:03        7023810      323095260              0              0
  00:00:c0:01:01:04        7023810      323095260              0              0
  00:00:c0:01:01:05        7023810      323095260              0              0
  00:00:c0:01:01:06        7023810      323095260              0              0
  00:00:c0:01:01:07        7023810      323095260              0              0
  00:00:c0:01:01:08        7023809      323095214              0              0
  00:00:c0:01:01:09        7023809      323095214              0              0
  00:00:c0:01:01:0a        7023809      323095214              0              0
  00:00:c0:01:01:0b        7023809      323095214              0              0
  00:00:c8:01:01:02       30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:03       30424784     1399540064       37448598     1722635508
  00:00:c8:01:01:04       30424716     1399536936       37448523     1722632058
  00:00:c8:01:01:05       30424789     1399540294       37448598     1722635508
  00:00:c8:01:01:06       30424788     1399540248       37448597     1722635462
  00:00:c8:01:01:07       30424783     1399540018       37448597     1722635462
  00:00:c8:01:01:08       30424783     1399540018       37448596     1722635416
  00:00:c8:01:01:09        8836796      406492616        8836795      406492570
  00:00:c8:01:01:0a       30424712     1399536752       37448521     1722631966
  00:00:c8:01:01:0b       30424715     1399536890       37448523     1722632058
  Number of MAC addresses : 21
```

show interfaces mac-database (All MAC Addresses on an Aggregated Ethernet Interface)

```
user@host> show interfaces mac-database ae4
Physical interface: ae4, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 588
  Description: Member links xe-0/2/0
  Link-level type: Ethernet, MTU: 9188, Speed: Unspecified, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:22:83:76:ff:c4, Hardware address: 00:22:83:76:ff:c4
  Last flapped   : 2013-12-18 04:33:22 PST (1w5d 22:23 ago)
  Input rate     : 62756384 bps (85266 pps)
  Output rate    : 62759472 bps (85272 pps)


  Logical interface ae4.0 (Index 334) (SNMP ifIndex 647)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    MAC address         Input frames    Input bytes  Output frames   Output bytes

    00:00:00:aa:00:02       23888711     2627758118            300          22200

    00:00:00:aa:00:03              0              0              0              0
```

```
    00:00:00:aa:00:04             0              0              0              0

    Number of MAC addresses : 3
```

show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address         Input frames     Input bytes  Output frames   Output bytes
  00:00:00:00:00:00              1              56              0              0
  00:00:c0:01:01:02        7023810       323095260              0              0
  00:00:c0:01:01:03        7023810       323095260              0              0
  00:00:c0:01:01:04        7023810       323095260              0              0
  00:00:c0:01:01:05        7023810       323095260              0              0
  00:00:c0:01:01:06        7023810       323095260              0              0
  00:00:c0:01:01:07        7023810       323095260              0              0
  00:00:c0:01:01:08        7023809       323095214              0              0
  00:00:c0:01:01:09        7023809       323095214              0              0
  00:00:c0:01:01:0a        7023809       323095214              0              0
  00:00:c0:01:01:0b        7023809       323095214              0              0
  00:00:c8:01:01:02       31016568      1426762128       38040381     1749857526
  00:00:c8:01:01:03       31016568      1426762128       38040382     1749857572
  00:00:c8:01:01:04       31016499      1426758954       38040306     1749854076
  00:00:c8:01:01:05       31016573      1426762358       38040381     1749857526
  00:00:c8:01:01:06       31016573      1426762358       38040381     1749857526
  00:00:c8:01:01:07       31016567      1426762082       38040380     1749857480
  00:00:c8:01:01:08       31016567      1426762082       38040379     1749857434
  00:00:c8:01:01:09        9428580       433714680        9428580      433714680
  00:00:c8:01:01:0a       31016496      1426758816       38040304     1749853984
  00:00:c8:01:01:0b       31016498      1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address  00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags      : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes      : 202324652
    Output bytes     : 202324560
    Input frames     : 4398362
    Output frames    : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```


show interfaces mc-ae

Syntax

show interfaces mc-ae
extensive
revertive-info
<id *identifier* unit *number*>

Release Information

Command introduced in Junos OS Release 9.6.
revertive-info statement introduced in Junos OS Release 13.3
extensive statement introduced in Junos OS Release 16.1R1

Description

On MX Series routers with multichassis aggregated Ethernet (**ae*X***) interfaces, displays information about the **ae*X*** interfaces.

Options

extensive—(Optional) Display extensive information for multichassis aggregated Ethernet interface.

revertive-info—(Optional) Display revertive mode information for multichassis aggregated Ethernet interface.

identifier—(Optional) Identifier of the multichassis aggregated Ethernet interface.

number—(Optional) Specify the logical interface by unit number.

Required Privilege Level

view

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*

List of Sample Output

show interfaces mc-ae on page 1991
show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers) on page 1992
show interfaces mc-ae revertive-info on page 1992
show interfaces mc-ae extensive on page 1992
show interfaces mc-ae extensive (MX Series Router after a configuration exchange error) on page 1993

Output Fields

Table 134 on page 1990 lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 134: show interfaces mc-ae Output Fields

Output Field Name	Field Description
Member Link	Identifiers of the configured multichassis link aggregate interfaces configured interfaces.
Local Status	Status of the local link: **active** or **standby**.

Table 134: show interfaces mc-ae Output Fields (continued)

Output Field Name	Field Description
Peer Status	Status of the peer link: **active** or **standby**.
Local State	Up or down state of the local device.
Peer State	Status of the local and peer links in an **active/active** bridge or VRRP over integrated routing and bridging (IRB) configuration on MX Series routers, including: **Logical Interface**—Aggregated Ethernet (AE) aggregate number and unit number. **Topology Type**—The bridge or VRRP topology type configured on the AE. **Local State**—Up or down state of the local device. **Peer State**—Up or down state of the peer device. **Peer Ip/ICL-PL/State**—Address, interface and state of the peer device.
Logical Interface	Identifier and unit of the multichassis aggregated Ethernet interface.
Core Facing Interface	Label: **pseudowire interface** or **Ethernet interface**.
ICL-PL	Label: **pseudowire interface** or **Ethernet interface**.
switchover mode	The configured switchover mode for the multichassis aggregated Ethernet interface: **revertive** or **non-revertive**.
switchover status	Status of the switchover if the **revert-time** statement is configured at the **[edit interfaces ae*x* mc-ae]** hierarchy level.
revert time	Revert time configured for the multichassis aggregated Ethernet interface.
switchover time remaining	Seconds left to trigger the switchover if the switchover is in progress.
Configuration Error Status	Reason for the configuration error.

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
```

```
 Logical Interface        : ae0.512
 Core Facing Interface : Label Ethernet Interface
 ICL-PL         : Label Ethernet Interface
```

show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers)

```
user@host# show interfaces mc-ae ge-0/0/0.0
 Member Link                  : ae0
 Current State Machine's State: active
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/ICL-PL/State     : 192.168.100.10 ge-0/0/0.0 up
```

show interfaces mc-ae revertive-info

```
user@host> show interfaces mc-ae revertive-info id 2
Member Link           : ae1
Current State Machine's State: mcae active state
Local Status                 : active
Local State                  : up
Peer Status                  : standby
Peer State                   : up
Switchover Mode              : Non Revertive
Switchover Status            : N/A
Revert Time                  : 1 Minutes
Switchover Remaining Time    : N/A
     Logical Interface        : ae1.1024
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : N/A
```

show interfaces mc-ae extensive

```
user@host> show interfaces mc-ae extensive
 Member Link                  : ae2
Current State Machine's State: mcae active state
Local Status                 : active
Local State                  : up
Peer Status                  : active
Peer State                   : up
     Logical Interface        : ae2.1
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.143.17 ae0.1 up

MCAE Configuration
     Redundancy Group         : 1
     MCAE ID                  : 2
     MCAE Mode                : active_active
     Status Control           : active
     Chassis ID               : 0
LACP Configuration
```

```
     System ID                : 00:00:00:00:00:02
     Admin Key                : 10
```

show interfaces mc-ae extensive (MX Series Router after a configuration exchange error)

```
user@host> show interfaces mc-ae extensive
 Member Link                    : ae2
Current State Machine's State: mcae config exchange error
Configuration Error Status   : same chassis-id
Local Status                 : active
Local State                  : up
Peer Status                  : Unknown
Peer State                   : Unknown
     Logical Interface        : ae2.1
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : 192.168.143.17 ae0.1 up

MCAE Configuration
     Redundancy Group         : 1
     MCAE ID                  : 2
     MCAE Mode                : active_active
     Status Control           : active
     Chassis ID               : 1
LACP Configuration
    System ID                : 00:00:00:00:00:02
     Admin Key                : 10
```

show interfaces transport pm

Syntax

show interfaces transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all |*interface-name*)

Release Information

Command introduced in Junos OS Release 14.2 on the PTX Series.
Command introduced in Junos OS Release 16.1 on the MX Series.

Description

Display diagnostic data, warnings, and alarms for transport performance monitoring interfaces.

Options

(all | optics | otn)—Display both optics and OTN information or either only optics or only OTN information.

(all | current | currentday | interval | previousday)—Display information for the current 15-minute interval, the current day, the ninety-six 15-minute intervals, and the previous day; information only for the current 15-minute interval; information only for the current 24 hours; information only for the ninety-six 15-minute intervals; information only for the previous day.

(all | *interface-name*)—Display information for all interfaces or only for the specified interface (for example, *et-fpc/pic/port*).

Required Privilege Level

view

Related Documentation

- clear interfaces transport pm on page 1462
- tca on page 974
- transport-monitoring on page 976

List of Sample Output

show interfaces transport pm on page 1996
show interfaces transport (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC) on page 1997
show interfaces transport pm (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC) on page 1997
show interfaces transport (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC) on page 1998
show interfaces transport pm optics (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC) on page 1998
show interfaces transport pm otn (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC) on page 1999

Output Fields

Table 135 on page 1995 lists the output fields for the **show interfaces transport pm optics** command. Fields are listed in the approximate order in which they appear.

Table 135: show interfaces transport pm Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Interval	The 15 minute interval for performance monitoring.
Suspect Flag	TRUE if the performance monitoring data for the interval appears to
Reason	Reason for setting the suspect flag.
COUNT	Measured value.
THRESHOLD	Threshold value set.
TCA-ENABLED	Threshold crossing alert. Set to TRUE if enabled.
TCA-RAISED	TRUE if enabled and the value crosses the threshold.
Near End PM	Near end threshold crossing defect trigger. For more information abo
Far End PM	Far end threshold crossing defect trigger. For more information abou
FEC PM	Forwarding equivalence class threshold crossing defect trigger. For m see tca.
BER PM	Bit error rate threshold crossing defect trigger. For more information a
CURRENT	Current value measured.
PM	Performance monitor.
MIN	Minimum value measured.
MAX	Maximum value measured.
AVG	Average value.
Lane Chromatic dispersion	Residual chromatic dispersion measured.
Lane differential group delay	Measured differential group delay.
q Value	Measured Quality factor value.
SNR	Signal to noise ratio.
Tx output power	The transmit laser output power.

Table 135: show interfaces transport pm Output Fields (continued)

Field Name	Field Description
Rx input power	The laser's received optical power.
Module temperature (Celsius)	The laser's temperature.
Tx Laser bias current (0.1mA)	Magnitude of the laser bias power setting current. The laser bias prov and modulates currents.
Rx Laser bias current (0.1mA)	Magnitude of the laser bias power setting current.
Carrier frequency offset (MHz)	Measured carrier frequency offset.

Sample Output

show interfaces transport pm

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False            Reason:None
    PM                  COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                0                800              No              No
  OTU-ES                 0                135              No              No
  OTU-SES                0                90               No              No
  OTU-UAS                427              90               No              No
  Far End        Suspect Flag:True             Reason:Unknown
    PM                  COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE                0                800              No              No
  OTU-ES                 0                135              No              No
  OTU-SES                0                90               No              No
  OTU-UAS                0                90               No              No
  Near End       Suspect Flag:False            Reason:None
    PM                  COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                0                800              No              No
  ODU-ES                 0                135              No              No
  ODU-SES                0                90               No              No
  ODU-UAS                427              90               No              No
  Far End        Suspect Flag:True             Reason:Unknown
    PM                  COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE                0                800              No              No
  ODU-ES                 0                135              No              No
  ODU-SES                0                90               No              No
  ODU-UAS                0                90               No              No
  FEC            Suspect Flag:False            Reason:None
    PM                  COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  FEC-CorrectedErr        2008544300       0                NA              NA
  FEC-UncorrectedWords    0                0                NA              NA
  BER            Suspect Flag:False            Reason:None
    PM                 MIN        MAX        AVG        THRESHOLD       TCA-ENABLED
```

```
     TCA-RAISED
  BER                 3.6e-5    5.8e-5    3.6e-5     10.0e-3       No
       Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current
  Suspect Flag:True             Reason:Object Disabled
    PM                             CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                             (MIN)
(MAX)    (MIN) (MAX)        (MIN) (MAX)
  Lane chromatic dispersion       0        0         0         0         0
 0         NA    NA          NA     NA
  Lane differential group delay   0        0         0         0         0
 0         NA    NA          NA     NA
  q Value                         120      120       120       120       0
 0         NA    NA          NA     NA
  SNR                             28       28        29        28        0
 0         NA    NA          NA     NA
  Tx output power(0.01dBm)        -5000    -5000     -5000     -5000     -300
 -100      No    No          No     No
  Rx output power(0.01dBm)        -3642    -3665     -3626     -3637     -1800
 -500      No    No          No     No
  Module temperature(Celsius)     46       46        46        46        -5
 75        No    No          No     No
  Tx laser bias current(0.1mA)    0        0         0         0         0
 0         NA    NA          NA     NA
  Rx laser bias current(0.1mA)    1270     1270      1270      1270      0
 0         NA    NA          NA     NA
  Carrier frequency offset(MHz)   -186     -186      -186      -186      -5000
 5000      No    No          No     No
```

show interfaces transport (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host > show interfaces transport et-3/0/0
Administrative State:       In Service
Operational    State:       Normal
```

show interfaces transport pm (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```
user@host > show interfaces transport pm otn current et-3/0/0
Physical interface: et-3/0/0, SNMP ifIndex 564
  23:30-current            Elapsed time:455 Seconds
  Near End       Suspect Flag:False             Reason:Not Applicable
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE               0                800              No              No
  OTU-ES                0                135              No              No
  OTU-SES               0                90               No              No
  OTU-UAS               0                90               No              No
  Far End        Suspect Flag:False             Reason:Not Applicable
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  OTU-BBE               0                800              No              No
  OTU-ES                0                135              No              No
  OTU-SES               0                90               No              No
  OTU-UAS               0                90               No              No
  Near End       Suspect Flag:False             Reason:Not Applicable
    PM                 COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED

  ODU-BBE               0                800              No              No
```

```
ODU-ES                  0               135             No              No
ODU-SES                 0               90              No              No
ODU-UAS                 0               90              No              No
Far End        Suspect Flag:False            Reason:Not Applicable
  PM                   COUNT           THRESHOLD      TCA-ENABLED     TCA-RAISED

ODU-BBE                 0               800             No              No
ODU-ES                  0               135             No              No
ODU-SES                 0               90              No              No
ODU-UAS                 0               90              No              No
FEC            Suspect Flag:False            Reason:Not Applicable
  PM                   COUNT           THRESHOLD      TCA-ENABLED     TCA-RAISED

FEC-CorrectedErr        30865849        0               NA              NA
FEC-UncorrectedWords    0               0               NA              NA
BER            Suspect Flag:False            Reason:Not Applicable
  PM                  MIN       MAX       AVG       THRESHOLD      TCA-ENABLED
   TCA-RAISED
BER                   4.0e-7    5.9e-7    5.1e-7     1.0e-2         No
     No
```

```
user@host > show interfaces transport pm optics current et-2/0/0
Physical interface: et-3/0/0, SNMP ifIndex 564
  23:30-current
  Suspect Flag:True               Reason:Not Applicable
    PM                              CURRENT   MIN       MAX       AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                             (MIN)
(MAX)    (MIN)  (MAX)       (MIN)  (MAX)
  Lane chromatic dispersion(ps/nm)  0        0         0          51         0
 0        NA     NA          NA     NA
  Lane differential group delay(ps) -13      13        0          11         0
 0        NA     NA          NA     NA
  q Value(0.1dB)                    0        -1        5          137        0
 0        NA     NA          NA     NA
  SNR(0.1dB)                        137      138       137        86         0
 0        NA     NA          NA     NA
  Tx output power(0.01dBm)          83       95        83         142        -300
 -100     No     No          No     No
  Rx input power(0.01dBm)           141      142       141        106        -1800
 -500     No     No          No     No
  Module temperature(Celsius)       106      109       106        -31        -5
 75       No     No          No     No
  Tx laser bias current(0.1mA)      -31      0         0          38         0
 0        NA     NA          NA     NA
  Rx laser bias current(0.1mA)      38       38        38         0          0
 0        NA     NA          NA     NA
  Carrier frequency offset(MHz)     0        0         2          0          -5000
 5000     No     No          No     No
```

show interfaces transport (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show interfaces transport et-8/0/0
Administrative State:       In Service
Operational    State:       Normal
```

show interfaces transport pm optics (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show interfaces transport pm optics current et-4/0/0
```



```
Physical interface: et-4/0/0, SNMP ifIndex 544
  02:45-current
  Suspect Flag:False            Reason:Not Applicable
    PM                             CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                             (MIN)
(MAX)     (MIN)  (MAX)        (MIN)  (MAX)
  Lane chromatic dispersion(ps/nm)  -6       -32       45        -1        0
 0        NA     NA          NA      NA
  Lane differential group delay(ps) 3        2         4         3         0
 0        NA     NA          NA      NA
  Lane Q2 factor(0.1dB)              154     154       155       154       0
 0        NA     NA          NA      NA
  SNR(0.1dB)                         167     164       171       165       0
 0        NA     NA          NA      NA
  Carrier frequency offset(MHz)     0        0         0         0         -3600
 3600      No     No          No      No
  Tx output power(0.01dBm)          0        0         0         0         -1100
 300       No     No          No      No
  Rx input total power(0.01dBm)     0        0         0         0         -3000
 300       No     No          No      No
  Module temperature(Celsius)        53      53        55        53        -5
 75        No     No          No      No
```

show interfaces transport pm otn (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host> show interfaces transport pm otn previousday et-4/0/0
Physical interface: et-4/0/0, SNMP ifIndex 544
  02:45-current
  Suspect Flag:False            Reason:Not Applicable
    PM                             CURRENT   MIN       MAX      AVG       THRESHOLD
      TCA-ENABLED       TCA-RAISED
                                                                             (MIN)
(MAX)     (MIN)  (MAX)        (MIN)  (MAX)
  Lane chromatic dispersion(ps/nm)  -6       -32       45        -1        0
 0        NA     NA          NA      NA
  Lane differential group delay(ps) 3        2         4         3         0
 0        NA     NA          NA      NA
  Lane Q2 factor(0.1dB)              154     154       155       154       0
 0        NA     NA          NA      NA
  SNR(0.1dB)                         167     164       171       165       0
 0        NA     NA          NA      NA
  Carrier frequency offset(MHz)     0        0         0         0         -3600
 3600      No     No          No      No
  Tx output power(0.01dBm)          0        0         0         0         -1100
 300       No     No          No      No
  Rx input total power(0.01dBm)     0        0         0         0         -3000
 300       No     No          No      No
  Module temperature(Celsius)        53      53        55        53        -5
 75        No     No          No      No
```


show l2-learning instance

Syntax **show l2-learning instance**

Release Information (MX Series routers only) Command introduced in Junos OS Release 8.4.

Description Display Layer 2 learning properties for all the configured routing instances.

Options This command has no options.

Required Privilege Level view

List of Sample Output show l2-learning instance on page 2001

Output Fields Table 136 on page 2000 describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

Table 136: show l2-learning instance Output Fields

Field Name	Field Description
Routing Instance	Name of routing instance.
Bridging Domain	Name of bridging domain. On MX Series routers you can use the **show l2-learning instance <extensive>** command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID.
Index	Number associated with the routing instance or bridging domain.
Logical System	Name of logical system or **Default** if no logical system is configured.
Routing instance flags	Status of Layer 2 learning properties for each routing instance: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after MAC address limit is reached. • **LH**—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.
MAC limit	Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.

Sample Output

show l2-learning instance

```
user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
                        AD -packet action drop, LH -mac limit hit)

Routing              Bridging     Index   Logical        Routing      MAC
Instance             Domain               System         flags        limit
__juniper_private1__                  1   Default                      5000
vs1                  vlan100          3   Default                      5120
vs1                  vlan200          4   Default                      5120
```

show l2-learning redundancy-groups

Syntax

show l2-learning redundancy-groups
logical-system [*system-name* | all]
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs

Release Information

Command introduced in Junos OS Release 13.2.
Support for logical systems added in Junos OS Release 14.1.
Command introduced in Junos OS Release 15.1R1 for EX Series switches

Description

(MX Series routers only) Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch or logical systems on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

Options

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

redundancy-group-id—(Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.

arp-statistics—(Optional) Count of ARP packets sent and received by the two MC-AE nodes.

nd-statistics—(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.

remote-macs —(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

Required Privilege Level

view

Related Documentation

- *Configuring Multichassis Link Aggregation on MX Series Routers*
- show interfaces mc-ae on page 1990
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation*
- *Configuring Multichassis Link Aggregation on EX Series Switches*

List of Sample Output

show l2-learning redundancy-groups arp-statistics on page 2004
show l2-learning redundancy-groups nd-statistics on page 2004
show l2-learning redundancy-groups remote-macs on page 2005

show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems) on page 2005
show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems) on page 2005
show l2-learning redundancy-groups group-id on page 2005
show l2-learning redundancy-groups logical-system on page 2005

Output Fields Output fields are listed in the approximate order in which they appear.

Table 137: show l2-learning redundancy-groups arp-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ARP Statistics Group ID	ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ARP Rx Count From Line	Total number of ARPs received from the Line.
ARP Tx Count To Peer	Total number of ARPs sent to the peer.
ARP Rx Count From Peer	Total number of ARPs received from the peer.
ARP Drop Count received from line	Total number of ARPs sent by the peer that were received.
ARP Drop Count received from peer	Total number of ARPs sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 138: show l2-learning redundancy-groups nd-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ND Statistics Group ID	Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ND Rx Count From Line	Total number of Neighbor Discovery packets received from the Line.
ND Tx Count To Peer	Total number of Neighbor Discovery packets sent to the peer.
ND Rx Count From Peer	Total number of Neighbor Discovery packets received from the peer.
ND Drop Count received from line	Total number of Neighbor Discovery packets sent by the peer that were received.

Table 138: show l2-learning redundancy-groups nd-statistics Output Fields (continued)

Field Name	Field Description
ND Drop Count received from peer	Total number of Neighbor Discovery packets sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 139: show l2-learning redundancy-groups remote-macs Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
Peer-Addr	IP address of the remote peer.
VLAN	Virtual LAN identifier associated with the redundancy group.
MAC	Hardware media access control address associated with the redundancy group.
MCAE-ID	ID number of the MC-AE used by the redundancy group.
Flags	Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected.
Status	Installation state: Installed or Not Installed.

Sample Output

show l2-learning redundancy-groups arp-statistics

```
user@host> show l2-learning redundancy-groups arp-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                 : 1
ARP Rx Count From Line                   : 52
ARP Tx Count To Peer                     : 15
ARP Rx Count From Peer                   : 39
ARP Install Count                        : 34
ARP Drop Count received from line        : 37
ARP Drop Count received from peer        : 5
```

show l2-learning redundancy-groups nd-statistics

```
user@host> show l2-learning redundancy-groups nd-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                 : 1
ND Rx Count From Line                   : 52
```



```
ND Tx Count To Peer                      : 15
ND Rx Count From Peer                    : 39
ND Install Count                         : 34
ND Drop Count received from line         : 37
ND Drop Count received from peer         : 5
```

show l2-learning redundancy-groups remote-macs

user@host> **show l2-learning redundancy-groups <redundancy-group-id> remote-macs**

```
  Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

Service-id  Peer-Addr    VLAN        MAC            MCAE-ID  Subunit  Opcode
Flags   Status
10          10.1.1.2      100   64:87:88:6a:df:f0      1        0       1
 0     Installed
```

show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems)

user@host> **show l2-learning redundancy-groups logical-system LS1 arp-statistics**

```
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                  : 1
ARP Rx Count From Line                    : 52
ARP Tx Count To Peer                      : 15
ARP Rx Count From Peer                    : 39
ARP Install Count                         : 34
ARP Drop Count received from line         : 37
ARP Drop Count received from peer         : 5
```

show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems)

user@host> **show l2-learning redundancy-groups logical-system LS1 nd-statistics**

```
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                  : 1
ND Rx Count From Line                    : 52
ND Tx Count To Peer                      : 15
ND Rx Count From Peer                    : 39
ND Install Count                         : 34
ND Drop Count received from line         : 37
ND Drop Count received from peer         : 5
```

show l2-learning redundancy-groups group-id

user@host> **show l2-learning redundancy-groups 1**

```
Redundancy Group ID : 1      Flags : Local Connect,Remote Connect
```

show l2-learning redundancy-groups logical-system

user@host> **show l2-learning redundancy-groups logical-system ls1**

```
Redundancy Group ID : 2      Flags : Local Connect,Remote Connect
```


show lacp interfaces

Syntax

show lacp interfaces
<*interface-name*>
extensive

Release Information

Command introduced in Junos OS Release 7.6.
extensive statement introduced in Junos OS Release 16.1R1
Command introduced in Junos OS Release 10.0 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.
Command introduced in Junos OS Release 14.2R3

Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

Options

none—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*number***
- Fast Ethernet—**fe-*fpc/pic/port***
- Gigabit Ethernet—**ge-*fpc/pic/port***
- 10 Gigabit Ethernet—**xe-*fpc/pic/port***

extensive—Display LACP information for the interface in detail.

NOTE: The show lacp interfaces command returns the following error message if your system is not configured in either active or passive LACP mode:

“Warning: lacp subsystem not running – not needed by configuration”

Required Privilege Level

view

Related Documentation

- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Configuring Link Aggregation*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*
- *Understanding Aggregated Ethernet Interfaces and LACP for MX Series*

- *Understanding Aggregated Ethernet Interfaces and LACP for Switches*
- *Junos OS Interfaces Fundamentals Configuration Guide*

List of Sample Output

show lacp interfaces (Aggregated Ethernet) on page 2010
show lacp interfaces (Gigabit Ethernet) on page 2010
show lacp interfaces (10 Gigabit Ethernet) on page 2011

Output Fields

Table 140 on page 2008 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 140: show lacp interfaces Output Fields

Field Name	Field Description	Level of Output
LACP State	For a child interface configured with the force-up statement, LACP state displays FUP along with the interface name.	All Levels
Aggregated interface	Aggregated interface value.	All Levels

LACP State	LACP state information for each aggregated interface:
	• **Role**—Role played by the interface. It can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **Exp**—Expired state. **Yes** indicates the actor or partner is in an expired state. **No** indicates the actor or partner is not in an expired state.
	• **Def**—Default. **Yes** indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. **No** indicates the operational partner information in use has been received in an LACP PDU.
	• **Dist**—Distribution of outgoing frames. **No** indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is **Yes**.
	• **Col**—Collection of incoming frames. **Yes** indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is **No**.
	• **Syn**—Synchronization. If the value is **Yes**, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is **No**, the link is not synchronized. It is currently not in the right aggregation.
	• **Aggr**—Ability of aggregation port to aggregate (**Yes**) or to operate only as an individual link (**No**).
	• **Timeout**—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (**Long Timeout** or **Short Timeout**).
	• **Activity**—Actor or partner's port activity. **Passive** indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is **Active**. **Active** indicates the port's preference to participate in the protocol regardless of the partner's control value.
		All Levels

Table 140: show lacp interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  ◦ **Current**—The state machine receives an LACP PDU and enters the **Current** state.
	  ◦ **Defaulted**—If no LACP PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state.
	  ◦ **Expired**—If no LACP PDU is received before the timer for the **Current** state expires once, the state machine enters the **Expired** state.
	  ◦ **Initialize**—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the **Initialize** state.
	  ◦ **LACP Disabled**—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to **LACP Disabled**. This state is similar to the **Defaulted** state, except that the port is forced to operate as an individual port.
	  ◦ **Port Disabled**—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the **Port Disabled** state.
	• **Transmit State**—Transmit state of state machine. One of the following values:
	  ◦ **Fast Periodic**—Periodic transmissions are enabled at a fast transmission rate.
	  ◦ **No Periodic**—Periodic transmissions are disabled.
	  ◦ **Periodic Timer**—Transitory state entered when the periodic timer expires.
	  ◦ **Slow Periodic**—Periodic transmissions are enabled at a slow transmission rate.
	• **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values:
	  ◦ **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.
	  ◦ **Collecting**—**Yes** indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. **No** indicates the receive function of this link is not enabled.
	  ◦ **Collecting Distributing**—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.
	  ◦ **Detached**—Process of detaching the port from the aggregator is in progress.
	  ◦ **Distributing**—**Yes** indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. **No** indicates the transmit function of this link is not enabled.
	  ◦ **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.
		All Levels

Table 140: show lacp interfaces Output Fields (continued)

Field Name	Field Description	Level of Output

LACP info	• **Role** can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **System priority**—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• **System identifier**—Actor or partner system ID, encoded as a MAC address.
	• **Port priority**—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• **Port number**—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.
	• **Port key**—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.
		Extensive

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
 LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/3      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/3    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/1                  Current   Fast periodic Collecting distributing
      ge-0/0/2                  Current   Fast periodic Collecting distributing
      ge-0/0/3                  Current   Fast periodic Collecting distributing
    LACP info:        Role     System             System       Port      Port
Port
                            priority         identifier   priority    number
 key
      ge-0/0/1       Actor        127   00:05:86:4e:b6:c0        127         1
   1
      ge-0/0/1     Partner        127   00:05:86:7e:d3:c0        127         1
   1
      ge-0/0/2       Actor        127   00:05:86:4e:b6:c0        127         2
   1
      ge-0/0/2     Partner        127   00:05:86:7e:d3:c0        127         2
   1
      ge-0/0/3       Actor        127   00:05:86:4e:b6:c0        127         3
   1
      ge-0/0/3     Partner        127   00:05:86:7e:d3:c0        127         3
   1
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/3/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
```

```
  ge-0/3/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
LACP Protocol:   Receive State    Transmit State          Mux State
  ge-0/3/0            Current     Fast periodic  Collecting distributing
```

show lacp interfaces (10 Gigabit Ethernet)

```
user@host> show lacp interfaces xe-1/0/2
Aggregated interface: ae0
    LACP State:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-1/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP Protocol:   Receive State    Transmit State          Mux State
      xe-1/0/2            Current     Fast periodic  Collecting distributing
```


show lldp

Syntax

show lldp
<detail>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display information about the Link Layer Discovery Protocol (LLDP).

Options

detail—(Optional) Display the detailed output level.

Required Privilege Level

view

List of Sample Output

show lldp on page 2014
show lldp detail on page 2014

Output Fields

Table 141 on page 2012 describes the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

Table 141: show lldp Output Fields

Field Name	Field Description
LLDP	Status of LLDP: **Enabled** or **Disabled**.
Advertisement interval	Value of the advertisement interval parameter.
Transmit delay	Value of the transmit delay parameter.
Hold timer	Value of the hold timer parameter.
Notification interval	Value of the notification interval parameter.
Config Trap Interval	Value of the configuration trap parameter.
Connection Hold timer	Value of the connection hold timer parameter.
Port ID TLV subtype	• *interface-name*—Indicates the interface name as the port information for the local device. • **locally-assigned**—Indicates that the sub-type for port ID TLV generation is locally assigned value of SNMP index of the interface. For more information about port ID TLV subtype, see port-id-subtype..

Table 141: show lldp Output Fields (continued)

Field Name	Field Description
Port Description TLV type	Following value used for port description TLV: • **interface-alias (ifAlias)**—Indicates that the *ifAlias* MIB object value is used to generate the port description TLV. • **interface-description (ifDescr)**—Indicates that the *ifDescr* MIB object value is used to generate the port description TLV. For more information about port description TLV type, see port-description-type.
Interface	Name of the interface for which LLDP configuration information is being reported For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs.
LLDP	LLDP operating state. The state can be Enabled or Disabled.
LLDP-MED	LLDP-MED operating state. The state can be Enabled or Disabled.
Power Negotiation	LLDP power negotiation operating state. The state can be Enabled or Disabled.
LLDP basic TLVs supported	List of basic LLDP TLVs supported by this device (**detail** only).
LLDP 802 TLVs supported	List of IEEE 802.1 LLDP TLVs supported by this device (**detail** only).

Sample Output

show lldp

```
user@host> show lldp
LLDP                        : Enabled
Advertisement interval      : 30 seconds
Transmit delay              : 2 seconds
Hold timer                  : 120 seconds
Notification interval       : 0 Second(s)
Config Trap Interval        : 0 seconds
Connection Hold timer       : 300 seconds
Port ID TLV subtype         : locally-assigned
Port Description TLV type : interface-description (ifDescr)

Interface      Parent Interface    LLDP       LLDP-MED       Power Negotiation
all            -                   Enabled
```

Sample Output

show lldp detail

```
user@host> show lldp detail
LLDP                        : Enabled
Advertisement interval      : 30 seconds
Transmit delay              : 2 seconds
Hold timer                  : 120 seconds
Notification interval       : 0 Second(s)
Config Trap Interval        : 0 seconds
Connection Hold timer       : 300 seconds
Port ID TLV subtype         : locally-assigned
Port Description TLV type : interface-description (ifDescr)

Interface      Parent Interface    LLDP       LLDP-MED     Power Negotiation
  Neighbor count
all            -                   Enabled
  2

Interface      Parent Interface    Vlan-id    Vlan-name
xe-0/0/0       -                   4080       vlan-4080
xe-0/0/1       -                   4080       vlan-4080

Basic Management TLVs supported:
End Of LLDPDU, Chassis ID, Port ID, Time To Live, Port Description, System Name,
 System Description, System Capabilities, Management Address

Organizationally Specific TLVs supported:
Port VLAN tag, VLAN Name, MAC/PHY Configuration/Status, Link Aggregation,Maximum
 Frame Size
```


show lldp local-information

Syntax show lldp local-information

Release Information Command introduced in Junos OS Release 9.6.

Description Display local Link Layer Discovery Protocol (LLDP) information.

Options This command has no options.

Required Privilege Level view

List of Sample Output
show lldp local-information(Management Information Address Subtype is IPv4) on page 2017
show lldp local-information(Management Information Address Subtype is IPv6) on page 2017

Output Fields Table 142 on page 2015 describes the output fields for the **show lldp local-information** command. Output fields are listed in the approximate order in which they appear.

Table 142: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	Information that follows pertains to the local system.
Chassis ID	List of chassis identifiers for local information.
System name	Local system name reported by LLDP.
System descr	Local system description reported by LLDP.
System Capabilities	Capabilities (such as **Bridge** or **Router**) that are **Supported** or **Enabled** by system on the interface.
Management Information	Listed by **Interface Name**, **Address Subtype** (such as **ipv4**, **ipv6**), **Address** (such as **192.168.168.229**, **1fd::1a10**), **Interface Number**, and **Interface Numbering Subtype**.
Interface Name	List of local interfaces. For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.
Parent Interface	Name of the **ae** interface to which the interface belongs
Interface ID	List of local interface identifiers.

Table 142: show lldp local-information Output Fields (continued)

Field Name	Field Description
Interface Description	List of local interface descriptions.
Status	List of interface conditions: **UP** or **DOWN**.

Sample Output

show lldp local-information(Management Information Address Subtype is IPv4)

```
user@host> show lldp local-information
LLDP Local Information details

Chassis ID   : 64:87:88:65:37:c0
System name  : apg-hp1
System descr : Juniper Networks, Inc. mx240 , version 14.1I20131231_0701_builder
 [builder] Build date: 2013-12-31 07:13:42 UTC

System Capabilities
    Supported       : Bridge Router
    Enabled         : Bridge Router

Management Information
    Interface Name  : Unknown
    Address Subtype : IPv4(1)
    Address         : 10.216.97.103
    Interface Number               : 1
    Interface Numbering Subtype    : ifIndex(2)

Interface name   Parent Interface   Interface ID   Interface description   Status
fxp0               -                 1              fxp0                     Up
me0                -                 33             me0                      Up
ge-2/0/0         ae0                1475           ge-2/0/0                 Up
ge-2/0/1         ae0                1476           ge-2/0/1                 Up
```

show lldp local-information(Management Information Address Subtype is IPv6)

```
user@host> show lldp local-information
LLDP Local Information details

Chassis ID   : ac:4b:c8:92:67:c0
System name  : apg-hp
System descr : Juniper Networks, Inc. mx240 , version 13.2-20131210.0 [builder]
               Build date: 2013-12-10 06:23:15 UTC

System Capabilities
    Supported       : Bridge Router
    Enabled         : Bridge Router

Management Information
    Interface Name  : fxp0
    Address Subtype : IPv6(2)
    Address         : 1fd::1a20
    Interface Number               : 1
    Interface Numbering Subtype    : ifIndex(2)

Interface name  Parent Interface   Interface ID   Interface description    Status
ge-1/2/4        -                  530            -                          Down
ge-1/2/5        -                  531            -                          Down
ge-1/2/2        -                  528            ge-1/2/2                    Up
ge-1/2/3        -                  529            ge-1/2/3                    Up
```


show lldp neighbors

Syntax

show lldp neighbors
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display information about LLDP neighbors.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .

Options

interface ***interface-name***—(Optional) Display the neighbor information about a particular physical interface.

NOTE: Starting with Junos OS Release 14.2, you can also display LLDP neighbor details for management interfaces, such as fxp or me, on MX Series routers.

Required Privilege Level

view

Related Documentation

- clear lldp neighbors on page 1463

List of Sample Output

show lldp neighbors on page 2021
show lldp neighbors interface ge-0/0/4 (Management Address is IPv4) on page 2021
show lldp neighbors interface ge-0/0/4 (Management Address is IPv6) on page 2022
show lldp neighbors (Management Ethernet Interfaces) on page 2023

Output Fields

Table 143 on page 2018 describes the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 143: show lldp neighbors Output Fields

Field Name	Field Description
LLDP Remote Devices Information	Information about remote devices.

Table 143: show lldp neighbors Output Fields (continued)

Field Name	Field Description
LocalInterface	List of local interfaces for which neighbor information is available.
ChassisId	List of chassis identifiers for neighbors.
PortInfo	List of port information gathered from neighbors. This could be the port identifier or port description.
SysName	List of system names gathered from neighbors.
LLDP Neighbor Information	Information about both local and neighbor systems on the interface (appears when the **interface** option is used).
Local Information	Information about local systems on the interface (appears when the **interface** option is used).
Neighbor Information	Information about both local and neighbor system on the interface (appears when the **interface** option is used).
Index	Local interface index (appears when the **interface** option is used).
Time Mark	Date and timestamp of information (appears when the **interface** option is used).
Time To Live	Number of seconds for which this information is valid (appears when the **interface** option is used).
Local Interface	Name of the local physical interface (appears when the **interface** option is used).
Parent Interface	Name of the **ae** interface to which the interface belongs
Local Port ID	Local port identifier (appears when the interface option is used).
Neighbor Information	Information about neighbor systems on the interface (appears when the **interface** option is used).
Chassis type	Type of chassis identifier supplied, such as **MAC address** (appears when the **interface** option is used).
Chassis ID	Chassis identifier of type listed (appears when the **interface** option is used).
Port type	Type of port identifier supplied, such as **local** (appears when the **interface** option is used).
Port ID	Port identifier of type listed (appears when the **interface** option is used).
Port description	Port description (appears when the **interface** option is used).

Table 143: show lldp neighbors Output Fields (continued)

Field Name	Field Description
System name	Name supplied by the system on the interface (appears when the **interface** option is used).
System Description	Description supplied by the system on the interface (appears when the **interface** option is used).
System Capabilities	Capabilities (such as **bridge** or **router**) that are **Supported** or **Enabled** by the system on the interface (appears when the **interface** option is used).
Management address	Details of the management address: **Address Type** (such as **ipv4** and **ipv6**), **Address** (such as **10.204.34.35**, **1fd::1a10**), **Interface Number**, **Interface Subtype**, and Organization Identifier (**OID**) (appears when the **interface** option is used).
Organization Info	One or more entries listing remote information by Organizationally Unique Identifier (**OUI**), **Subtype**, **Index**, and **Info** (appears when the **interface** option is used).

Sample Output

show lldp neighbors

```
user@host> show lldp neighbors

Local Interface  Parent Interface  Chassis Id         Port info     System Name
ge-2/0/0         ae0               ac:4b:c8:92:67:c0   528          apg-hp
ge-2/0/1         ae0               ac:4b:c8:92:67:c0   529          apg-hp
```

Sample Output

show lldp neighbors interface ge-0/0/4 (Management Address is IPv4)

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 2 Time to live: 120 Time mark: Tue Dec 31 11:47:46 2013 Age: 15 secs
Local Interface    : ge-2/0/1
Parent Interface   : ae0
Local Port ID      : 1476
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : ac:4b:c8:92:67:c0
Port type          : Locally assigned
Port ID            : 529
Port description   : ge-1/2/3
System name        : apg-hp

System Description : Juniper Networks, Inc. mx240 , version 14.1-20131222.0
[builder] Build date: 2013-12-22 09:13:26 UTC


System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv4(1)
        Address           : 10.216.98.57
        Interface Number  : 1
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
 Capability (0x1d), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status (0x3), Aggregation Port ID (1694498816)
       Index    : 2

Organization Info
```

```
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1518)
       Index    : 3
```

show lldp neighbors interface ge-0/0/4 (Management Address is IPv6)

```
user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 1 Time to live: 120 Time mark: Thu Dec 12 07:19:45 2013 Age: 28 secs
Local Interface    : ge-1/2/2
Parent Interface   : -
Local Port ID      : 528
Ageout Count      : 0

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : 64:87:88:65:37:c0
Port type          : Locally assigned
Port ID            : 1475
Port description   : ge-2/0/0
System name        : apg-hp1

System Description : Juniper Networks, Inc. mx240 , version 11.4R10 Build date:
2013-10-24 10:10:02 UTC


System capabilities
        Supported: Bridge Router
        Enabled  : Bridge Router

Management address
        Address Type      : IPv6(2)
        Address           : 1fd::1a10
        Interface Number  : 1
        Interface Subtype : ifIndex(2)
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : MAC/PHY Configuration/Status (1)
       Info     : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
 Capability (0x5), MAU Type (0x0)
       Index    : 1

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Link Aggregation (3)
       Info     : Aggregation Status (0x1), Aggregation Port ID (0)
       Index    : 2

Organization Info
       OUI      : IEEE 802.3 Private (0x00120f)
       Subtype  : Maximum Frame Size (4)
       Info     : MTU Size (1518)
       Index    : 3

Organization Info
       OUI      : Ethernet Bridged (0x0080c2)
```



```
Subtype   : VLAN Name (3)
Info      : VLAN ID (100), VLAN Name (vlan-100)
Index     : 4
```

show lldp neighbors (Management Ethernet Interfaces)

user@host> **show lldp neighbors**

```
Local Interface    Parent Interface    Chassis Id          Port info
System Name
fxp0               -                   78:fe:3d:ee:4e:00   151
x2-sw35
xe-0/0/0           -                   a8:d0:e5:50:26:c0   512
sitara
xe-0/0/1           -                   a8:d0:e5:50:26:c0   513
sitara
```

show lldp remote-global-statistics

Syntax show lldp remote-global-statistics

Release Information Command introduced in Junos OS Release 9.6.

Description Display remote Link Layer Discovery Protocol (LLDP) global statistics.

Options This command has no options.

Required Privilege Level view

List of Sample Output show lldp remote-global-statistics on page 2025

Output Fields Table 144 on page 2024 describes the output fields for the **show lldp remote-global-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 144: show lldp remote-global-statistics Output Fields

Field Name	Field Description
LLDP Remote Database Table Counters	Information about remote database table counters.
LastchangeTime	Time elapsed between LLDP agent startup and the last change to the remote database table information.
Inserts	Number of insertions made in the remote database table.
Deletes	Number of deletions made in the remote database table.
Drops	Number of LLDP frames dropped from the remote database table because of errors.
Ageouts	Number of remote database table entries that have aged out of the table.

Sample Output

show lldp remote-global-statistics

```
user@host> show lldp remote-global-statistics
user@host> show lldp remote-global-statistics
   LLDP Remote Database Table Counters
   LastchangeTime          Inserts    Deletes    Drops    Ageouts
   00:00:76 (76 sec)       192        0          0        0
```


show lldp statistics

Syntax

show lldp statistics
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

Display information about Link Layer Discovery Protocol (LLDP) statistics.

Options

interface *interface-name*—(Optional) Display the statistics about a particular physical interface.

NOTE: Starting with Junos OS Release 14.2, you can also display LLDP statistical details for management interfaces, such as fxp or me, on MX Series routers.

Required Privilege Level

view

Related Documentation

- clear lldp statistics on page 1464

List of Sample Output

show lldp statistics on page 2028
show lldp statistics interface ge-0/1/1 on page 2028

Output Fields

Table 145 on page 2026 describes the output fields for the **show lldp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 145: show lldp statistics Output Fields

Field Name	Field Description
Interface	Interface name. For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*. For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports* .
Received	Number of LLDP frames received on this interface.
Transmitted	Number of LLDP frames sent on this interface.

Table 145: show lldp statistics Output Fields (continued)

Field Name	Field Description
Unknown-TLVs	Number of LLDP frames with unsupported content received on this interface.
With-Errors	Number of LLDP frames with errors received on this interface.
Discarded	Number of LLDP frames received on this interface that were discarded because of problems.
Transmitted	Total number of LLDP frames that were transmitted on an interface.
Untransmitted	Total number of LLDP frames that were untransmitted on an interface.

Sample Output

show lldp statistics

```
user@host> show lldp statistics

Interface     Parent Interface  Received  Unknown TLVs  With Errors
xe-3/0/0.0    ae31.0            1564      0             0
xe-3/0/1.0    ae31.0            1564      0             0
xe-3/0/2.0    ae31.0            1565      0             0
xe-3/0/3.0    ae31.0            1566      0             0
xe-3/0/4.0    ae31.0            1598      0             0
xe-3/0/5.0    ae31.0            1598      0             0
xe-3/0/6.0    ae31.0            1596      0             0
xe-3/0/7.0    ae31.0            1597      0             0
xe-5/0/6.0    -                 0         0             0
xe-5/0/7.0    -                 0         0             0


Discarded TLVs  Transmitted  Untransmitted
0               3044         1
0               3044         1
0               3044         1
0               3044         1
0               3075         1
0               3075         1
0               3075         1
0               3075         1
0               17312        0
0               17312        0
```

Sample Output

show lldp statistics interface ge-0/1/1

```
user@host> show lldp statistics interface ge-0/1/1
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
--------- -------- ----------- ------------ ----------- ---------
ge-0/1/1  544      540         0            0           0
```

show oam ethernet connectivity-fault-management delay-statistics

Syntax

show oam ethernet connectivity-fault-management delay-statistics
<count *entry-count*>
<local-mep *mep-id*>
<maintenance-association *ma-name*>
<maintenance-domain *md-name*>
<remote-mep *remote-mep-id*>

Release Information

Command introduced in Junos OS Release 9.5.
Command introduced in Junos OS Release 11.4 for EX Series switches.

Description

On MX Series routers with Ethernet interfaces on Dense Port Concentrators (DPCs), display ETH-DM delay statistics.

On EX Series switches, display delay measurement results.

Options

count *entry-count*—(Optional) Number of entries to display from the statistics table. The range of values is **1** through **100**. The default value is **100** entries.

local-mep *mep-id*—(Optional) Numeric identifier of the local MEP. On MX Series routers, the range of values is **1** through **8192**. On EX Series switches, the range of values is **1** through **8191**.

maintenance-association *ma-name*—Name of an existing CFM maintenance association.

maintenance-domain *md-name*—Name of an existing connectivity fault management (CFM) maintenance domain.

remote-mep *remote-mep-id*—(Optional) Numeric identifier of the remote MEP. On MX Series routers, the range of values is **1** through **8192**. On EX Series switches, the range of values is **1** through **8191**.

Required Privilege Level

view

Related Documentation

- clear oam ethernet connectivity-fault-management statistics on page 1469
- *clear oam ethernet connectivity-fault-management delay-statistics*
- show oam ethernet connectivity-fault-management interfaces on page 2037
- show oam ethernet connectivity-fault-management mep-database on page 2055
- show oam ethernet connectivity-fault-management mep-statistics on page 2066

List of Sample Output

show oam ethernet connectivity-fault-management delay-statistics on page 2031
show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 2031

Output Fields Table 146 on page 2030 lists the output fields for the **show oam ethernet connectivity-fault-management delay-statistics** command and the **show oam ethernet connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 146: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.
Two-way delay (usec)	For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.

Sample Output

show oam ethernet connectivity-fault-management delay-statistics

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec

  Remote MEP identifier: 102
  Remote MAC address: 00:04:55:63:39:5a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        29            58
      2        23            59
      3        27            56
      4        29            62
      5        33            68
    Average one-way delay          : 28 usec
    Average one-way delay variation: 3 usec
    Best case one-way delay        : 23 usec
    Worst case one-way delay       : 33 usec
    Average two-way delay          : 60 usec
    Average two-way delay variation: 3 usec
    Best case two-way delay        : 56 usec
    Worst case two-way delay       : 68 usec
```

show oam ethernet connectivity-fault-management delay-statistics remote-mep

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
```

```
  1       259            519
  2       273            550
  3       287            571
  4       299            610
  5       313            650
Average one-way delay         : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay       : 259 usec
Worst case one-way delay      : 313 usec
Average two-way delay         : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay       : 519 usec
Worst case two-way delay      : 650 usec
```


show oam ethernet connectivity-fault-management forwarding-state

Syntax

show oam ethernet connectivity-fault-management forwarding-state
interface ***interface-name*** **| instance** ***instance-name***
<brief | detail | extensive>

Release Information

Command introduced in Junos OS Release 8.4.

Description

On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.

Options

interface ***interface-name***—Display forwarding state information for the specified Ethernet interface only.

instance ***instance-name***—Display forwarding state information for the specified forwarding instance only.

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

Required Privilege Level

view

List of Sample Output

show oam ethernet connectivity-fault-management forwarding-state instance on page 2034
show oam ethernet connectivity-fault-management forwarding-state interface on page 2034
show oam ethernet connectivity-fault-management forwarding-state interface detail on page 2035
show oam ethernet connectivity-fault-management forwarding-state interfaceinterface-name on page 2036

Output Fields

Table 147 on page 2033 lists the output fields for the **show oam ethernet connectivity-fault-management forwarding-state** command. Output fields are listed in the approximate order in which they appear.

Table 147: show oam ethernet connectivity-fault-management forwarding-state Output Fields

Field Name	Field Description	Level of Output
Interface name	Interface identifier.	All levels
Link (Status)	Local link status.	All levels
Filter action	Filter action for messages at the level.	All levels
Next hop type	Next-hop type.	All levels
Next index	Next-hop index number.	**brief**

Table 147: show oam ethernet connectivity-fault-management forwarding-state Output Fields (continued)

Field Name	Field Description	Level of Output
Level	Maintenance domain (MD) level.	**detail**
Direction	MEP direction configured.	none
Instance name	Forwarding instance name.	All levels
CEs	Number of customer edge (CE) interfaces.	All levels
VEs	Number of VPN endpoint (VE) interfaces.	All levels

Sample Output

show oam ethernet connectivity-fault-management forwarding-state instance

```
user@host> show oam ethernet connectivity–fault-management forwarding-state instance
Instance name: __+bd1__
CEs: 3
VEs: 0
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7                       Drop              none
```

show oam ethernet connectivity-fault-management forwarding-state interface

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-3/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
```



```
  7        down           Receive          none

Interface name: xe-0/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action    Nexthop      Nexthop
                                           type         index
  0                       Drop             none
  1                       Drop             none
  2                       Drop             none
  3                       Drop             none
  4                       Drop             none
  5                       Drop             none
  6                       Drop             none
  7        down           Receive          none
```

show oam ethernet connectivity-fault-management forwarding-state interface detail

user@host> **show oam ethernet connectivity-fault-management forwarding-state interface detail**

```
Interface name: ge-3/0/0.0
Instance name: __+bd1__

  Level: 0
  Filter action: Drop
  Nexthop type: none

  Level: 1
  Filter action: Drop
  Nexthop type: none

  Level: 2
  Filter action: Drop
  Nexthop type: none

  Level: 3
  Filter action: Drop
  Nexthop type: none

  Level: 4
  Filter action: Drop
  Nexthop type: none

  Level: 5
  Filter action: Drop
  Nexthop type: none

  Level: 6
  Filter action: Drop
  Nexthop type: none

  Level: 7
  Direction: down
  Filter action: Receive
  Nexthop type: none

Interface name: xe-0/0/0.0
```



```
Instance name: __+bd1__

  Level: 0
  Filter action: Drop
  Nexthop type: none

  Level: 1
  Filter action: Drop
  Nexthop type: none

  ...
```

show oam ethernet connectivity-fault-management forwarding-state interface interface-name

user@host> **show oam ethernet connectivity-fault-management forwarding-state interface interface-name ge-3/0/0/0.0**

```
Interface name: ge-3/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action      Nexthop      Nexthop
                                             type         index
  0                       Drop               none
  1                       Drop               none
  2                       Drop               none
  3                       Drop               none
  4                       Drop               none
  5                       Drop               none
  6                       Drop               none
  7        down           Receive            none
```

show oam ethernet connectivity-fault-management interfaces

Syntax

show oam ethernet connectivity-fault-management interfaces
<*ethernet-interface-name*>
<level *md-level*>
<brief | detail | extensive>

Release Information

Command introduced in Junos OS Release 8.4.
Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.
Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for ACX Series and MX Series routers.

Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.

In addition, for Ethernet interfaces on MX Series routers , also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts when **detail** or **extensive** mode is specified.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

brief | detail | extensive—(Optional) Specified level of output.

ethernet-interface-name—(Optional) CFM information only for CFM entities attached to the specified Ethernet interface.

level *md-level*—(Optional) CFM information for CFM identities enclosed within a maintenance domain of the specified level.

Required Privilege Level

view

Related Documentation

- clear oam ethernet connectivity-fault-management statistics on page 1469
- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- show oam ethernet connectivity-fault-management mep-database on page 2055
- show oam ethernet connectivity-fault-management mep-statistics on page 2066
- Ethernet Alarm Indication Signal (ETH-AIS) Function Overview on page 805

List of Sample Output

show oam ethernet connectivity-fault-management interfaces on page 2042
show oam ethernet connectivity-fault-management interfaces detail on page 2043
show oam ethernet connectivity-fault-management interfaces detail (One-Way ETH-DM) on page 2044

show oam ethernet connectivity-fault-management interfaces detail (Connection Protection TLV Configured) on page 2044
show oam ethernet connectivity-fault-management interfaces extensive on page 2045
show oam ethernet connectivity-fault-management interfaces level on page 2046
show oam ethernet connectivity-fault-management interfaces (trunk ports) on page 2046

Output Fields

Table 148 on page 2038 lists the output fields for the **show oam ethernet connectivity-fault-management interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 148: show oam ethernet connectivity-fault-management interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Interface identifier.	All levels
Interface status	Local interface status.	All levels
Link status	Local link status. **Up**, **down**, or **oam-down**.	All levels
Maintenance domain name	Maintenance domain name.	**detail extensive**
Format (Maintenance domain)	Maintenance domain name format configured.	**detail extensive**
Level	Maintenance domain level configured.	All levels
Maintenance association name	Maintenance association name.	**detail extensive**
Format (Maintenance association)	Maintenance association name format configured.	**detail extensive**
Continuity-check status	Continuity-check status.	**detail extensive**
Ethernet-ais status	Status of alarm indication signal (AIS). **active** or **in-active**.	**detail extensive**
Interval	Continuity-check message interval.	**detail extensive**
Loss-threshold	Lost continuity-check message threshold.	**detail extensive**
Interface status TLV	Status of the interface status TLV, if configured on the MEP interface: **none**, **up**, **down**, **testing**, **unknown**, **dormant**, **notPresent**, **lowerLayerDown**	**detail extensive**
Port status TLV	Status of the port status TLV, if configured on the MEP interface: **none**, **no**, **yes**	**detail extensive**

Table 148: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Connection Protection TLV	Status of the connection protection TLV if configured on the MEP interface: **no**, **yes** If **yes**, then the transmitted connection protection TLV is decoded and the following three fields are displayed: **Prefer me**, **Protection in use**, **FRR Flag**	**detail extensive**
Prefer me	If set to **yes**, the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to the remote side. Its value can be **yes** or **no**.	**detail extensive**
Protection in use	Used for protection decision coordination. Its value is set to **yes** if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path. Its value can be **yes** or **no**.	**detail extensive**
FRR Flag	LSR/LER forwarding the CCM Frame into a bypass tunnel is set. Its value can be **yes** or **no**.	**detail extensive**
MEP identifier	Maintenance association end point (MEP) identifier.	All levels
Neighbors	Number of MEP neighbors.	All levels
Direction	MEP direction configured.	**detail extensive**
MAC address	MAC address configured for the MEP.	**detail extensive**
MEP status	Indicates the status of the connectivity fault management (CFM) protocol running on the MEP: **Running**, **inactive**, **disabled**, or **unsupported**.	**detail extensive**
Remote MEP not receiving CCM	Whether the remote MEP is not receiving connectivity check messages (CCMs).	**detail extensive**
Erroneous CCM received	Whether erroneous CCMs have been received.	**detail extensive**
Cross-connect CCM received	Whether cross-connect CCMs have been received.	**detail extensive**
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.	**detail extensive**
Some remote MEP's MAC in error state	Indicates whether the remote MEP’s MAC is in error state.	**detail extensive**
Alarm Indication Signal	Indicates whether the AIS is triggered or is cleared.	**detail extensive**

Table 148: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CCMs sent	Number of CCMs transmitted.	**detail extensive**
CCMs received out of sequence	Number of CCMs received out of sequence.	**detail extensive**
LBMs sent	Number of loopback request messages (LBMs) sent.	**detail extensive**
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.	**detail extensive**
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.	**detail extensive**
LBRs received with corrupted data	Number of LBRs received that were corrupted.	**detail extensive**
LBRs sent	Number of LBRs transmitted.	**detail extensive**
LTMs sent	Linktrace messages (LTMs) transmitted.	**detail extensive**
LTMs received	Linktrace messages received.	**detail extensive**
LTRs sent	Linktrace responses (LTRs) transmitted.	**detail extensive**
LTRs received	Linktrace responses received.	**detail extensive**
Sequence number of next LTM request	Sequence number of next LTM request to be transmitted.	**detail extensive**
1DMs sent	If the interface is attached to an initiator MEP for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.	**detail extensive**
Valid 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of valid 1DM frames received. For all other cases, this field displays **0**.	**detail extensive**
Invalid 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of invalid 1DM frames received. For all other cases, this field displays **0**.	**detail extensive**
Out of sync 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.	**detail extensive**

Table 148: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
DMMs sent	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.	**detail extensive**
Valid DMMs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid two-way delay measurement request packets received.	**detail extensive**
Invalid DMMs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid two-way delay measurement request packets received.	**detail extensive**
DMRs sent	If the interface is attached to a responder MEP for a two-way ETH-DM session: Number of delay measurement reply (DMR) frames sent. For all other cases, this field displays **0**.	**detail extensive**
Valid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays **0**.	**detail extensive**
Invalid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays **0**.	**detail extensive**
LMM sent	If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.	**detail extensive**
Valid LMM received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.	**detail extensive**
Invalid LMM received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid loss measurement request packets received.	**detail extensive**
LMR sent	If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.	**detail extensive**
Valid LMR received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid LMR frames received.	**detail extensive**
Invalid LMR received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.	**detail extensive**
Valid AIS frames transmitted	Number of valid AIS frames transmitted to the peer MEPs.	**detail extensive**
Valid AIS frames received	Number of valid AIS frames received from the peer MEPs.	**detail extensive**

Table 148: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
SLM sent	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session.	**detail extensive**
Valid SLM received	If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.	**detail extensive**
Invalid SLM received	If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.	**detail extensive**
SLR sent	If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.	**detail extensive**
Valid SLR received	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.	**detail extensive**
Invalid SLR received	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.	**detail extensive**
Remote MEP count	Number of remote MEPs.	**extensive**
Identifier (remote MEP)	MEP identifier of the remote MEP.	**extensive**
MAC address (remote MEP)	MAC address of the remote MEP.	**extensive**
State (remote MEP)	State of the remote MEP.	**extensive**
Interface (remote MEP)	Interface of the remote MEP.	**extensive**

Sample Output

show oam ethernet connectivity-fault-management interfaces

```
user@host> show oam ethernet connectivity-fault-management interfaces
Interface       Link      Status           Level       MEP          Neighbors
                                                       Identifier
ge-1/1/0.0      Up        Active           0           2            1
ge-1/1/0.1      Up        Active           0           2            1
ge-1/1/0.10     Up        Active           0           2            1
ge-1/1/0.100    Up        Active           0           2            1
ge-1/1/0.101    Up        Active           0           2            1
ge-1/1/0.102    Up        Active           0           2            1
```

```
ge-1/1/0.103    Up        Active          0           2            1
ge-1/1/0.104    Up        Active          0           2            1
ge-1/1/0.105    Up        Active          0           2            1
ge-1/1/0.106    Up        Active          0           2            1

...
```

show oam ethernet connectivity-fault-management interfaces detail

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Maintenance domain name: md0, Format: string, Level: 5
  Maintenance association name: ma1, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
   MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                : no
      Erroneous CCM received                      : yes
      Cross-connect CCM received                  : no
      RDI sent by some MEP                        : yes
      Alarm Indication Signal                     : yes
    Statistics:
      CCMs sent                                   : 76
      CCMs received out of sequence               : 0
      LBMs sent                                   : 0
      Valid in-order LBRs received                : 0
      Valid out-of-order LBRs received            : 0
      LBRs received with corrupted data           : 0
      LBRs sent                                   : 0
      LTMs sent                                   : 0
      LTMs received                               : 0
      LTRs sent                                   : 0
      LTRs received                               : 0
      Sequence number of next LTM request         : 1320235363

      1DMs sent                                   : 0
      Valid 1DMs received                         : 0
      Invalid 1DMs received                       : 0
      DMMs sent                                   : 0
      DMRs sent                                   : 0
      Valid DMRs received                         : 0
      Invalid DMRs received                       : 0
      LMM sent                                    : 10
      Valid LMM received                          : 20
      Invalid LMM received                        : 0
      LMR sent                                    : 20
      Valid LMR received                          : 10
      Invalid LMR received                        : 0
      Valid AIS frames transmitted                : 0
      Valid AIS frames received                    : 0
      SLM sent                                    : 10
      Valid SLM received                          : 20
      Invalid SLM received                        : 0
      SLR sent                                    : 20
      Valid SLR received                          : 10
      Invalid SLR received                        : 0
    Remote MEP count: 2
      Identifier     MAC address        State    Interface
```



```
2001      00:90:69:0b:7f:71       ok     ge-5/2/9.0
4001      00:90:69:0b:09:c5       ok     ge-5/2/9.0
```

show oam ethernet connectivity-fault-management interfaces detail (One-Way ETH-DM)

```
user@host show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
     MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                  : no
      Erroneous CCM received                        : no
      Cross-connect CCM received                    : no
      RDI sent by some MEP                          : no
    Statistics:
      CCMs sent                                     : 1590
      CCMs received out of sequence                 : 0
      LBMs sent                                     : 0
      Valid in-order LBRs received                  : 0
      Valid out-of-order LBRs received              : 0
      LBRs received with corrupted data             : 0
      LBRs sent                                     : 0
      LTMs sent                                     : 0
      LTMs received                                 : 0
      LTRs sent                                     : 0
      LTRs received                                 : 0
      Sequence number of next LTM request           : 1542035464

      1DMs sent                                     : 10
      Valid 1DMs received                           : 0
      Invalid 1DMs received                         : 0
      DMMs sent                                     : 0
      DMRs sent                                     : 0
      Valid DMRs received                           : 0
      Invalid DMRs received                         : 0
    Remote MEP count: 1
      Identifier     MAC address         State    Interface
        201      00:90:69:0a:43:94       ok     ge-0/2/5.0
```

show oam ethernet connectivity-fault-management interfaces detail (Connection Protection TLV Configured)

```
user@hostshow oam ethernet connectivity-fault-management interfaces detail

Interface name: xe-6/2/0.0 , Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
 Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: none, Port status TLV: none
    Connection Protection TLV: yes
      Prefer me: no, Protection in use: no, FRR Flag: no
    MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                  : no
```

```
      Erroneous CCM received                     : no
      Cross-connect CCM received                 : no
      RDI sent by some MEP                       : no
      Some remote MEP's MAC in error state       : no
    Statistics:
      CCMs sent                                  : 225
      CCMs received out of sequence              : 0
      LBMs sent                                  : 0
      Valid in-order LBRs received               : 0
      Valid out-of-order LBRs received           : 0
      LBRs received with corrupted data          : 0
      LBRs sent                                  : 0
      LTMs sent                                  : 0
      LTMs received                              : 0
      LTRs sent                                  : 0
      LTRs received                              : 0
      Sequence number of next LTM request        : 1244305646

      1DMs sent                                  : 0
      Valid 1DMs received                        : 0
      Invalid 1DMs received                      : 0
      Out of sync 1DMs received                  : 0
      DMMs sent                                  : 0
      Valid DMMs received                        : 0
      Invalid DMMs received                      : 0
      DMRs sent                                  : 0
      Valid DMRs received                        : 0
      Invalid DMRs received                      : 0
      LMMs sent                                  : 0
      Valid LMMs received                        : 0
      Invalid LMMs received                      : 0
      LMRs sent                                  : 0
      Valid LMRs received                        : 0
      Invalid LMRs received                      : 0
    Remote MEP count: 1
      Identifier    MAC address        State    Interface
          2     00:90:69:7f:e4:30
```

show oam ethernet connectivity-fault-management interfaces extensive

```
user@host> show oam ethernet connectivity-fault-management interfaces extensive
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Maintenance domain name: md0, Format: string, Level: 5
    Maintenance association name: ma1, Format: string
    Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
    Interface status TLV: none, Port status TLV: none
    Connection Protection TLV: no
 MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
    MEP status: running
    Defects:
      Remote MEP not receiving CCM               : no
      Erroneous CCM received                     : yes
      Cross-connect CCM received                 : no
      RDI sent by some MEP                       : yes

      Alarm Indication Signal                    : yes
    Statistics:
      CCMs sent                                  : 76
      CCMs received out of sequence              : 0
```

```
      LBMs sent                                : 0
      Valid in-order LBRs received             : 0
      Valid out-of-order LBRs received         : 0
      LBRs received with corrupted data        : 0
      LBRs sent                                : 0
      LTMs sent                                : 0
      LTMs received                            : 0
      LTRs sent                                : 0
      LTRs received                            : 0
      Sequence number of next LTM request      : 1645032434

      1DMs sent                                : 0
      Valid 1DMs received                      : 0
      Invalid 1DMs received                    : 0
      DMMs sent                                : 0
      DMRs sent                                : 0
      Valid DMRs received                      : 0
      Invalid DMRs received                    : 0

      Valid AIS frames transmitted             : 0
      Valid AIS frames received                 : 0
      SLM sent                                 : 10
      Valid SLM received                       : 20
      Invalid SLM received                     : 0
      SLR sent                                 : 20
      Valid SLR received                       : 10
      Invalid SLR received                     : 0
    Remote MEP count: 2
      Identifier     MAC address        State    Interface
       2001      00:90:69:0b:7f:71       ok    ge-5/2/9.0
       4001      00:90:69:0b:09:c5       ok    ge-5/2/9.0
```

show oam ethernet connectivity-fault-management interfaces level

```
user@host> show oam ethernet connectivity-fault-management interfaces level 7
Interface        Link        Status            Level       MEP          Neighbors
                                                           Identifier
ge-3/0/0.0       Up          Active            7           201          0
xe-0/0/0.0       Up          Active            7           203          1
```

show oam ethernet connectivity-fault-management interfaces (trunk ports)

```
user@host> show oam ethernet connectivity-fault-management interfaces

Interface                    Link      Status       Level  MEP          Neighbors
                                                           Identifier
ge-4/0/1.0, vlan 100         Up        Active       5      100          0
ge-10/3/10.4091, vlan 4091   Down      Inactive     4      400          0
ge-4/0/0.0                   Up        Active       6      200          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0

Interface                    Link      Status       Level  MEP          Neighbors
                                                           Identifier
ge-4/0/0.0                   Up        Active       6      200          0
```

user@host> **show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vlan 100**

```
Interface                    Link      Status       Level  MEP         Neighbors
                                                           Identifier
ge-4/0/1.0, vlan 100         Up        Active       5      100         0
```

user@host> **show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091 vlan 4091**

```
Interface                    Link      Status       Level  MEP         Neighbors
                                                           Identifier
ge-10/3/10.4091, vlan 4091   Down      Inactive     4      400         0
```


show oam ethernet connectivity-fault-management linktrace path-database

Syntax

show oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-association *ma-name* maintenance-domain *md-name*

Release Information

Command introduced in Junos OS Release 9.0.

Description

On M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.

Options

mac-address—Display connectivity fault management path database information for the specified MAC address of the remote host.

maintenance-association *ma-name*—Display connectivity fault management path database information for the specified maintenance association.

maintenance-domain *md-name*—Display connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level

view

List of Sample Output

show oam ethernet connectivity-fault-management linktrace path-database on page 2049
show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 2049

Output Fields

Table 149 on page 2048 lists the output fields for the **show oam ethernet connectivity-fault-management linktrace path-database** command. Output fields are listed in the approximate order in which they appear.

Table 149: show oam ethernet connectivity-fault-management linktrace path-database Output Fields

Field Name	Field Description
Linktrace to	MAC address of the 802.1ag node to which the linktrace message is targeted.
Interface	Interface used by the local MEP to send the linktrace message (LTM).
Maintenance Domain	Maintenance domain identifier specified in the traceroute command.
Maintenance Association	Maintenance association identifier specified in the traceroute command.
Level	Maintenance domain level configured for the maintenance domain.
Local Mep	MEP identifier of the local MEP originating the linktrace.

Table 149: show oam ethernet connectivity-fault-management linktrace path-database Output Fields (continued)

Field Name	Field Description
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message (LTM). The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag node responding to the LTM or the source MAC address of the LTR.
Next hop MAC address	MAC address of the egress interface of the node to which the LTM is forwarded or the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming linktrace responses (LTR), with a previously sent LTM.

Sample Output

show oam ethernet connectivity-fault-management linktrace path-database

```
user@host> show oam ethernet connectivity-fault-management linktrace path-database
maintenance-domain MD1 maintenance-association MA1 00:01:02:03:04:05
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

        Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100001
       1        63       00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
       2        62       00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
       3        61       00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
       4        60       00:01:02:03:04:05         00:00:00:00:00:00
```

show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands)

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:01:02:03:04:05
 Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100002
       1        63       00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
       2        62       00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
       3        61       00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
       4        60       00:01:02:03:04:05         00:00:00:00:00:00
```



```
Transaction Identifier:100003
1       63       00:00:aa:aa:aa:aa          00:00:ab:ab:ab:ab
2       62       00:00:bb:bb:bb:bb          00:00:bc:bc:bc:bc
3       61       00:00:cc:cc:cc:cc          00:00:cd:cd:cd:cd
4       60       00:01:02:03:04:05          00:00:00:00:00:00
```


show oam ethernet connectivity-fault-management loss-statistics

Syntax

show oam ethernet connectivity-fault-management loss-statistics
maintenance-domain ***md-name***
maintenance-association ***ma-name***
<count ***entry-count*****>**
<local-mep ***local-mep-id*****>**
<remote-mep ***remote-mep-id*****>**

Release Information

Command introduced in Junos OS Release 11.1.

Description

On MX Series and ACX series routers with Ethernet interfaces, display ETH-LM statistics for on-demand mode only.

Options

maintenance-domain ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

count ***entry-count***—(Optional) Number of entries to display from the statistics table. The range of values is from 1 through 100. The default value is 100.

local-mep ***local-mep-id***—(Optional) Numeric identifier of the local MEP. The range of values is from 1 through 8191.

remote-mep ***remote-mep-id***—(Optional) Numeric identifier of the remote MEP. The range of values is from 1 through 8191.

Required Privilege Level

view

Related Documentation

- show oam ethernet connectivity-fault-management mep-statistics on page 2066

Output Fields

Table 150 on page 2051 lists the output fields for the **show oam ethernet connectivity-fault-management loss-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 150: show oam ethernet connectivity-fault-management loss-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
Remote MEP identifier	Numeric identifier of the remote MEP.

Table 150: show oam ethernet connectivity-fault-management loss-statistics Output Fields (continued)

Output Field Name	Field Description
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-LM entry in the CFM database.
Near-end frame loss	Count of frame loss associated with ingress data frames.
Far-end frame loss	Count of frame loss associated with egress data frames.
Near-end loss ratio	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the ingress interface.
Far-end loss ratio	Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the egress interface.
Average near-end frame loss	Average frame loss measured in this session associated with ingress data frames.
Average near-end loss ratio	Average frame loss ratio measured in this session associated with ingress data frames.
Average far-end frame loss	Average frame loss measured in this session associated with egress data frames.
Average far-end loss ratio	Average frame loss ratio measured in this session associated with egress data frames.
Near-end best case loss	Lowest frame loss measured in this session associated with ingress data frames.
Near-end best case loss ratio	Lowest frame loss ratio measured in this session associated with ingress data frames.
Near-end worst case loss	Highest frame loss measured in this session associated with ingress data frames.
Near-end worst case loss ratio	Highest frame loss ratio measured in this session associated with ingress data frames.
Far-end best case frame loss	Lowest frame loss measured in this session associated with egress data frames.
Far-end best case loss ratio	Lowest frame loss ratio measured in this session associated with egress data frames.
Far-end worst case loss	Highest frame loss measured in this session associated with egress data frames.
Far-end worst case loss ratio	Highest frame loss ratio measured in this session associated with egress data frames.

show oam ethernet connectivity fault management loss statistics

```
user@host>. show oam ethernet connectivity fault management loss statistics
maintenance-domain md maintenance-association ma
MEP identifier: 1, MAC address: 64:87:88:f9:7d:1b
Remote MEP count: 1

  Remote MAC address: 64:87:88:6a:da:94

    LM client session-id:4843
    CIR Loss measurement statistics:
    Index  Near-end      Far-end      Near-end       Far-end       Near-end
Far-end
           Frame loss    Total tx     Total rx       Frame loss    Total tx
Total rx
             (CIR)         (CIR)        (CIR)          (CIR)         (CIR)
  (CIR)
       1        0            245          245          0             244
      244
       2        0            488          488          0             489
      489
       3        0            732          732          0             733
      733
       4        0            977          977          0             976
      976
    EIR Loss measurement statistics:
    Index  Near-end      Far-end      Near-end       Far-end      Near-end
Far-end
           Frame loss    Total tx     Total rx       Frame loss    Total tx
Total rx
             (EIR)         (EIR)        (EIR)         (EIR)         (EIR)
 (EIR)
       1        0            272          272          0             273
      273
       2        0            546          546          0             545
      545
       3        0            820          820          0             819
      819
       4        0            1092         1092         0             1093
      1093
    Total far-end Tx (CIR)             : 977
    Total near-end Rx (CIR)            : 977
    Total near-end loss(CIR)           : 0
    Total near-end loss ratio(CIR)     : 0.00000%
    Total near-end Tx (CIR)            : 976
    Total far-end Rx (CIR)             : 976
    Total far-end loss(CIR)            : 0
    Total far-end loss ratio(CIR)      : 0.00000%
    Average near-end loss(CIR)         : 0.00000
    Average near-end loss ratio(CIR)   : 0.00000%
    Average far-end loss(CIR)          : 0.00000
    Average far-end loss ratio(CIR)    : 0.00000%
    Near-end best case loss(CIR)       : 0
    Near-end best case loss ratio(CIR) : 0.00000%
    Near-end worst case loss(CIR)      : 0
    Near-end worst case loss ratio(CIR): 0.00000%
    Far-end best case loss(CIR)        : 0
    Far-end best case loss ratio(CIR)  : 0.00000%
    Far-end worst case loss(CIR)       : 0
    Far-end worst case loss ratio(CIR) : 0.00000%
```



```
Total far-end Tx (EIR)             : 1092
Total near-end Rx (EIR)            : 1092
Total near-end loss(EIR)           : 0
Total near-end loss ratio(EIR)     : 0.00000%
Total near-end Tx (EIR)            : 1093
Total far-end Rx (EIR)             : 1093
Total far-end loss(EIR)            : 0
Total far-end loss ratio(EIR)      : 0.00000%
Average near-end loss(EIR)         : 0.00000
Average near-end loss ratio(EIR)   : 0.00000%
Average far-end loss(EIR)          : 0.00000
Average far-end loss ratio(EIR)    : 0.00000%
Near-end best case loss(EIR)       : 0
Near-end best case loss ratio(EIR) : 0.00000%
Near-end worst case loss(EIR)      : 0
Near-end worst case loss ratio(EIR): 0.00000%
Far-end best case loss(EIR)        : 0
Far-end best case loss ratio(EIR)  : 0.00000%
Far-end worst case loss(EIR)       : 0
Far-end worst case loss ratio(EIR) : 0.00000%
```


show oam ethernet connectivity-fault-management mep-database

Syntax

show oam ethernet connectivity-fault-management mep-database
maintenance-domain ***domain-name***
maintenance-association ***ma-name***
<local-mep ***local-mep-id*****>**
<remote-mep ***remote-mep-id*****>**

Release Information

Command introduced in Junos OS Release 8.4.
Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.
Support for ITU-T Y.1731 synthetic frame loss measurement added in Junos OS Release 13.2 for MX Series routers.

Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, M120, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

In addition, on M120, M320, and MX series routers, also display port status TLV, interface status TLV, and action profile information.

In addition, for Ethernet interfaces on MX Series routers , also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

maintenance-association ***ma-name***—Name of the maintenance association.

maintenance-domain ***domain-name***—Name of the maintenance domain.

local-mep-id—(Optional) Numeric identifier of local MEP.

remote-mep-id—(Optional) Numeric identifier of the remote MEP.

Required Privilege Level

view

Related Documentation

- clear oam ethernet connectivity-fault-management statistics on page 1469
- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- show oam ethernet connectivity-fault-management interfaces on page 2037
- show oam ethernet connectivity-fault-management mep-statistics on page 2066

List of Sample Output

show oam ethernet connectivity-fault-management mep-database on page 2060
show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM) on page 2061

show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 2062
show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 2062
show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured) on page 2062
show oam ethernet connectivity-fault-management mep-database on page 2063
show oam ethernet connectivity-fault-management mep-database (enhanced continuity measurement) on page 2064

Output Fields Table 151 on page 2056 lists the output fields for the **show oam ethernet connectivity-fault-management mep-database** command. Output fields are listed in the approximate order in which they appear.

Table 151: show oam ethernet connectivity-fault-management mep-database Output Fields

Field Name	Field Description
Maintenance domain name	Maintenance domain name.
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.
Format (Maintenance association)	Maintenance association name format configured.
Continuity-check status	Continuity-check status.
Interval	Continuity-check message interval.
Loss-threshold	Lost continuity-check message threshold.
Connection Protection TLV	Status of the connection protection TLV, if configured on the MEP interface: **no**, **yes** If **yes**, then the transmitted connection protection TLV is decoded and the following three fields are displayed: **Prefer me**, **Protection in use**, **FRR Flag**
Prefer me	If set to **yes**, the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to remote side. Its value can be **yes** or **no**.
Protection in use	Used for protection decision coordination. Its value is set to **yes** if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path. Its value can be **yes** or **no**.

Table 151: show oam ethernet connectivity-fault-management mep-database Output Fields (continued)

Field Name	Field Description
FRR Flag	LSR/LER forwarding the CCM Frame into a bypass tunnel is set. Its value can be **yes** or **no**.
MEP identifier	Maintenance association end point (MEP) identifier.
Direction	MEP direction configured.
MAC address	MAC address configured for the MEP.
Auto-discovery	Whether automatic discovery is enabled or disabled.
Priority	Priority used for CCMs and linktrace messages transmitted by the MEP.
Interface name	Interface identifier.
Interface status	Local interface status.
Link status	Local link status.
Remote MEP not receiving CCM	Whether the remote MEP is not receiving CCMs.
Erroneous CCM received	Whether erroneous CCMs have been received.
Cross-connect CCM received	Whether cross-connect CCMs have been received.
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
1DMs sent	If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.

Table 151: show oam ethernet connectivity-fault-management mep-database Output Fields (continued)

Field Name	Field Description
Valid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received. For all other cases, this field displays **0**.
Invalid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received. For all other cases, this field displays **0**.
Out of sync 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.
DMMs sent	If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.
Valid DMMs received	If the MEP is an initiator for a two-way ETH-DM session: Number of valid two-way delay measurement packets received.
Invalid DMMs received	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid two-way delay measurement packets received.
DMRs sent	If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent. For all other cases, this field displays **0**.
Valid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays **0**.
Invalid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays **0**.
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.
LBRs received with corrupted data	Number of LBRs received that were corrupted.
LBRs sent	Number of LBRs transmitted.
LTMs sent	Linktrace messages (LTMs) transmitted.
LTMs received	Linktrace messages received.
LTRs sent	Linktrace responses (LTRs) transmitted.
LTRs received	Linktrace responses received.

Table 151: show oam ethernet connectivity-fault-management mep-database Output Fields (continued)

Field Name	Field Description
Sequence number of next LTM request	Sequence number of the next linktrace message request to be transmitted.
LMM sent	If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.
Valid LMM received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.
Invalid LMM received	If the interface is attached to an initiator MEP for a ETH LM session: Number of invalid loss measurement request packets received.
LMR sent	If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.
Valid LMR received	If the interface is attached to an initiator MEP for a ETH LM session: Number of valid LMR frames received.
Invalid LMR received	If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.
SLM sent	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session.
Valid SLM received	If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.
Invalid SLM received	If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.
SLR sent	If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.
Valid SLR received	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.
Invalid SLR received	If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.
Remote MEP identifier	MEP identifier of the remote MEP.
State (remote MEP)	State of the remote MEP: **idle**, **start**, **ok**, or **failed**.
MAC address	MAC address of the remote MEP.
Type	Whether the remote MEP MAC address was learned using automatic discovery or configured.

Table 151: show oam ethernet connectivity-fault-management mep-database Output Fields (continued)

Field Name	Field Description
Interface	Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.
Last flapped	Date, time, and how long ago the remote MEP interface went from down to up. The format is **Last flapped:** ***year-month-day hours*:*minutes*:*seconds timezone*** **(*hours*:*minutes*:*seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.
Remote defect indication	Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.
Port status TLV	• In the Maintenance domain section, displays the last transmitted port status TLV value. • In the Remote MEP section, displays the last value of port status TLV received from the remote MEP. In the Action profile section, displays, the last occurred event **port-status-tlv blocked** event. This event occurred due to the reception of **blocked** value in the port status TLV from remote MEP.
Interface status TLV	• In the Maintenance domain section, displays the last transmitted interface status TLV value. • In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP. In the Action profile section, if displays, the last occurred event interface-status-tlv event (either **lower-layer-down** or **down**). This event occurred due to the reception of either lower or **down** value in the interface status TLV from remote MEP.
Action profile	Name of the action profile occurrence associated with a remote MEP.
Last event	When an action profile occurs, displays the last event that triggered it.
Last event cleared	When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.
Action	Action taken and the corresponding time of the action occurrence.

Sample Output

show oam ethernet connectivity-fault-management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
 Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: no  Interface name: ge-0/0/1.0, Interface status:
Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
```



```
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 1476
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 85
    Valid in-order LBRs received                  : 78
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 1
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 1
    Sequence number of next LTM request           : 1
    1DMs sent                                     : 0
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      100     00:19:e2:b2:81:4b       ok  vt-0/1/10.1049088
```

show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM)

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md6 maintenance-domain ma6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
  Statistics:
    CCMs sent                                     : 1590
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 10
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    DMMs sent                                     : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
```



```
  Remote MEP count: 1
    Identifier     MAC address        State    Interface
      201     00:90:69:0a:43:94       ok     ge-0/2/5.0
```

show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200 remote-mep 100**

```
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 100, State: ok
    MAC address: 00:19:e2:b2:81:4b, Type: Learned
    Interface: vt-0/1/10.1049088
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain md5 maintenance-association ma5 remote-mep 200**

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 200, State: ok
    MAC address: 00:05:85:73:96:1f, Type: Configured
    Interface: ge-1/0/8.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: lower-layer-down
    Action profile: juniper
      Last event: Interface-status-tlv lower-layer-down
      Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
```

show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured)

user@host>**show oam ethernet connectivity-fault-management mep-database maintenance-domain md5 maintenance-association ma5**

If connection-protection is not enabled on down MEPs, but connection-protection TLV is used, MX always sets the protection-in-use flag in connection-protection tlv, while CCMs are sent out. During reversion, this is an indicator to the receiver that protect-path is in use, otherwise the peer (receiver) assumes working is active and reversion does not work as expected. Setting this bit does not affect protection-switching/traffic-loss.

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-6/2/0.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                  : no
    Erroneous CCM received                        : no
    Cross-connect CCM received                    : no
    RDI sent by some MEP                          : no
    Some remote MEP's MAC in error state          : no
  Statistics:
    CCMs sent                                     : 251
    CCMs received out of sequence                 : 0
    LBMs sent                                     : 0
    Valid in-order LBRs received                  : 0
    Valid out-of-order LBRs received              : 0
    LBRs received with corrupted data             : 0
    LBRs sent                                     : 0
    LTMs sent                                     : 0
    LTMs received                                 : 0
    LTRs sent                                     : 0
    LTRs received                                 : 0
    Sequence number of next LTM request           : 0
    1DMs sent                                     : 0
    Valid 1DMs received                           : 0
    Invalid 1DMs received                         : 0
    Out of sync 1DMs received                     : 0
    DMMs sent                                     : 0
    Valid DMMs received                           : 0
    Invalid DMMs received                         : 0
    DMRs sent                                     : 0
    Valid DMRs received                           : 0
    Invalid DMRs received                         : 0
    LMMs sent                                     : 0
    Valid LMMs received                           : 0
    Invalid LMMs received                         : 0
    LMRs sent                                     : 0
    Valid LMRs received                           : 0
    Invalid LMRs received                         : 0
  Remote MEP count: 1
   Identifier    MAC address        State    Interface
        2     00:90:69:7f:e4:30
```

show oam ethernet connectivity-fault-management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:14:f6:b6:01:fe
Auto-discovery: enabled, Priority: 0
```

```
Interface name: ge-1/0/0.0, Interface status: Active, Link status: Up

Defects:
Remote MEP not receiving CCM         : no
Erroneous CCM received               : no
Cross-connect CCM received           : no
RDI sent by some MEP                 : no

Statistics:
CCMs sent                            : 328703
CCMs received out of sequence        : 0
LBMs sent                            : 85
Valid in-order LBRs received         : 78
Valid out-of-order LBRs received     : 0
LBRs received with corrupted data    : 0
LBRs sent                            : 0
LTMs sent                            : 0
LTMs received                        : 0
LTRs sent                            : 0
LTRs received                        : 0
Sequence number of next LTM request  : 0
1DMs sent                            : 10
Valid 1DMs received                  : 10
Invalid 1DMs received                : 0
DMMs sent                            : 20
DMRs sent                            : 0
Valid DMRs received                  : 10
Invalid DMRs received                : 0
LMM sent                             : 10
Valid LMM received                   : 20
Invalid LMM received                 : 0
LMR sent                             : 20
Valid LMR received                   : 10
Invalid LMR received                 : 0
SLM sent                                      : 10
Valid SLM received                            : 20
Invalid SLM received                          : 0
SLR sent                                      : 20
Valid SLR received                            : 10
Invalid SLR received                          : 0

Remote MEP count                     : 1

Identifier     MAC address         State     Interface
    2       00:12:1e:fb:ea:7d       ok        ge-1/0/0.0
```

show oam ethernet connectivity-fault-management mep-database (enhanced continuity measurement)

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001**

```
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address  : 00:19:e2:b0:74:00, Type: Learned
```



```
Interface    : ge-2/0/0.0
Last flapped  : Never
+ Continuity  : 91%, Admin-enable duration: 2100sec, Oper-down duration: 100sec
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
```

show oam ethernet connectivity-fault-management mep-statistics

Syntax

show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain ***md-name***
maintenance-association ***ma-name***
<mep ***mep-id*****>**
<remote-mep ***remote-mep-id*****>**
<count ***entry-count*****>**

Release Information

Command introduced in Junos OS Release 9.5.
Command introduced in Junos OS Release 11.4 for EX Series switches.
Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for MX Series routers.

Description

On MX Series and ACX Series routers and EX Series switches with Ethernet interfaces, display ETH-DM statistics and ETH-DM frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

maintenance-domain ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

mep ***mep-id***—(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.

remote-mep ***remote-mep-id***—(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.

count ***entry-count***—(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.

Required Privilege Level

view

Related Documentation

- clear oam ethernet connectivity-fault-management statistics on page 1469
- show oam ethernet connectivity-fault-management delay-statistics on page 2029
- show oam ethernet connectivity-fault-management interfaces on page 2037
- show oam ethernet connectivity-fault-management mep-database on page 2055

List of Sample Output

show oam ethernet connectivity-fault-management mep-statistics (CIR counters only) on page 2069
show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled) on page 2070

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only) on page 2071
show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled) on page 2073
show oam ethernet connectivity-fault-management mep-statistics on page 2074
show oam ethernet connectivity-fault-management mep-statistics remote-mep on page 2075
show oam ethernet connectivity-fault-management mep-statistics local-mep remote-mep on page 2076

Output Fields

Table 152 on page 2067 lists the output fields for the **show oam ethernet connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 152: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
CCMs sent	Number of continuity check messages (CCMs) sent.
CCMs received	Number of continuity check messages (CCMs) received for a specific remote MEP and maintenance association.
CCMs received out of sequence	Number of continuity check messages (CCMs) received that were not in sequence.
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.
Two-way delay (usec)	For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.
Average one-way delay	Average one-way frame delay for the statistics displayed.

Table 152: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields (continued)

Output Field Name	Field Description
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLM packets received	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
SLR packets sent	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFCl at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t	Value of the local counter RxFCl at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault-management mep-statistics (CIR counters only)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count                   : 1
CCMs sent                          : 6550
CCMs received out of sequence      : 0
LBMs sent                          : 0
Valid in-order LBRs received       : 0
Valid out-of-order LBRs received   : 0
LBRs received with corrupted data  : 0
LBRs sent                          : 0
LTMs sent                          : 0
LTMs received                      : 0
LTRs sent                          : 0
LTRs received                      : 0
Sequence number of next LTM request : 0
1DMs sent                          : 5
Valid 1DMs received                : 0
Invalid 1DMs received              : 0
DMMs sent                          : 5
DMRs sent                          : 0
Valid DMRs received                : 5
Invalid DMRs received              : 0
LMM sent                           : 5
Valid LMM received                 : 5
Invalid LMM received               : 0
LMR sent                           : 0
Valid LMR received                 : 5
Invalid LMR received               : 0
Remote MEP identifier              : 101
Remote MAC address                 : 00:05:85:73:39:4a

Delay measurement statistics:
Index        One-way delay        Two-way delay
                (usec)               (usec)
  1              259                 519
  2              273                 550
  3              287                 571
  4              299                 610
  5              313                 650

Average one-way delay                : 286 usec
Average one-way delay variation      : 62 usec
Best case one-way delay              : 259 usec
Average two-way delay                : 580 usec
Average two-way delay variation      : 26 usec
Best case two-way delay              : 519 usec
Worst case two-way delay             : 650 usec

Loss measurement statistics:
Index      Near-end     Far-end       Near-end     Far-end
           Frame loss   Frame loss    Frame loss   Frame loss
             (CIR)        (CIR)         (EIR)        (EIR)
  1            9            9
  2            3            5
```

```
  3           7           5
  4           9           6
  5           3           6

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)     : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)      : 6.2%
Near-end best case loss (CIR)         : 3
Near-end best case loss ratio (CIR)   : 3%
Near-end worst case loss (CIR)        : 9
Near-end worst case loss ratio (CIR)  : 9%
Far-end best case loss (CIR)          : 5
Far-end best case loss ratio (CIR)    : 5%
Far-end worst case loss (CIR)         : 9
Far-end worst case loss ratio (CIR)   : 9%
```

show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled)

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count                      : 1
CCMs sent                             : 6550
CCMs received out of sequence         : 0
LBMs sent                             : 0
Valid in-order LBRs received          : 0
Valid out-of-order LBRs received      : 0
LBRs received with corrupted data     : 0
LBRs sent                             : 0
LTMs sent                             : 0
LTMs received                         : 0
LTRs sent                             : 0
LTRs received                         : 0
Sequence number of next LTM request   : 0
1DMs sent                             : 5
Valid 1DMs received                   : 0
Invalid 1DMs received                 : 0
DMMs sent                             : 5
DMRs sent                             : 0
Valid DMRs received                   : 5
Invalid DMRs received                 : 0
LMM sent                              : 5
alid LMM received                     : 5
Invalid LMM received                  : 0
LMR sent                              : 0
Valid LMR received                    : 5
Invalid LMR received                  : 0
Remote MEP identifier                 : 101
Remote MAC address                    : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay          Two-way delay
              (usec)                 (usec)
  1            259                    519
  2            273                   550
  3            287                   571
  4            299                   610
  5            313                   650
```

```
Average one-way delay                  : 286 usec
Average one-way delay variation        : 62 usec
Best case one-way delay                : 259 usec
Average two-way delay                  : 580 usec
Average two-way delay variation        : 26 usec
Best case two-way delay                : 519 usec
Worst case two-way delay               : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
           (CIR)         (CIR)        (EIR)         (EIR)
  1          9             9            2             4
  2          3             5            4             6
  3          7             5            0             2
  4          9             6            8             2
  5          3             6            6             4

Average near-end loss (CIR)               : 6.2
Average near-end loss ratio (CIR)         : 6.2%
Average far-end loss (CIR)                : 6.2
Average far-end loss ratio (CIR)          : 6.2%
Near-end best case loss (CIR)             : 3
Near-end best case loss ratio (CIR)       : 3%
Near-end worst case loss (CIR)            : 9
Near-end worst case loss ratio (CIR)      : 9%
Far-end best case loss (CIR)              : 5
Far-end best case loss ratio (CIR)        : 5%
Far-end worst case loss (CIR)             : 9
Far-end worst case loss ratio (CIR)       : 9%
Average near-end loss (EIR)               : 4
Average near-end loss ratio (EIR)         : 4%
Average far-end loss (EIR)                : 3.4
Average far-end loss ratio (EIR)          : 3.4%
Near-end best case loss (EIR)             : 0
Near-end best case loss ratio (EIR)       : 0%
Near-end worst case loss (EIR)            : 8
Near-end worst case loss ratio (EIR)      : 8%
Far-end best case loss (EIR)              : 2
Far-end best case loss ratio (EIR)        : 2%
Far-end worst case loss (EIR)             : 6
Far-end worst case loss ratio (EIR)       : 6%
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only)

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3 remote-mep 101**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                             : 7762
CCMs received out of sequence         : 0
LBMs sent                             : 0
Valid in-order LBRs received          : 0
Valid out-of-order LBRs received      : 0
LBRs received with corrupted data     : 0
LBRs sent                             : 0
LTMs sent                             : 0
LTMs received                         : 0
LTRs sent                             : 0
LTRs received                         : 0
```



```
Sequence number of next LTM request   : 0
1DMs sent                            : 5
Valid 1DMs received                  : 0
Invalid 1DMs received                : 0
DMMs sent                            : 5
DMRs sent                            : 0
Valid DMRs received                  : 5
Invalid DMRs received                : 0
LMM sent                             : 5
Valid LMM received                   : 5
Invalid LMM received                 : 0
LMR sent                             : 0
Valid LMR received                   : 5
Invalid LMR received                 : 0
Remote MEP identifier                : 101
Remote MAC address                   : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay        Two-way delay
               (usec)               (usec)
  1             259                  519
  2             273                  550
  3             287                  571
  4             299                  610
  5             313                  650

Average one-way delay                 : 286 usec
Average one-way delay variation       : 62 usec
Best case one-way delay               : 259 usec
Average two-way delay                 : 580 usec
Average two-way delay variation       : 26 usec
Best case two-way delay               : 519 usec
Worst case two-way delay              : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)        (CIR)         (EIR)        (EIR)
  1           9            9
  2           3            5
  3           7            5
  4           9            6
  5           3            6
Average near-end loss (CIR)              : 6.2
Average near-end loss ratio (CIR)        : 6.2%
Average far-end loss (CIR)               : 6.2
Average far-end loss ratio (CIR)         : 6.2%
Near-end best case loss (CIR)            : 3
Near-end best case loss ratio (CIR)      : 3%
Near-end worst case loss (CIR)           : 9
Near-end worst case loss ratio (CIR)     : 9%
Far-end best case loss (CIR)             : 5
Far-end best case loss ratio (CIR)       : 5%
Far-end worst case loss (CIR)            : 9
Far-end worst case loss ratio (CIR)      : 9%
Average near-end loss (EIR)              : 4
Average near-end loss ratio (EIR)        : 4%
Average far-end loss (EIR)               : 3.4
Average far-end loss ratio (EIR)         : 3.4%
Near-end best case loss (EIR)            : 0
Near-end best case loss ratio (EIR)      : 0%
```

```
Near-end worst case loss (EIR)           : 8
Near-end worst case loss ratio (EIR)     : 8%
Far-end best case loss (EIR)             : 2
Far-end best case loss ratio (EIR)       : 2%
Far-end worst case loss (EIR)            : 6
Far-end worst case loss ratio (EIR)      : 6%
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled)

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3 remote-mep 101**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                           :7762
CCMs received out of sequence       : 0
LBMs sent                           : 0
Valid in-order LBRs received        : 0
Valid out-of-order LBRs received    : 0
LBRs received with corrupted data   : 0
LBRs sent                           : 0
LTMs sent                           : 0
LTMs received                       : 0
LTRs sent                           : 0
LTRs received                       : 0
Sequence number of next LTM request : 0
1DMs sent                           : 5
Valid 1DMs received                 : 0
Invalid 1DMs received               : 0
DMMs sent                           : 5
DMRs sent                           : 0
Valid DMRs received                 : 5
Invalid DMRs received               : 0
LMM sent                            : 5
Valid LMM received                  : 5
Invalid LMM received                : 0
LMR sent                            : 0
Valid LMR received                  : 5
Invalid LMR received                : 0
Remote MEP identifier               : 101
Remote MAC address                  : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
  4             299                   610
  5             313                   650

Average one-way delay                : 286 usec
Average one-way delay variation      : 62 usec
Best case one-way delay              : 259 usec
Average two-way delay                : 580 usec
Average two-way delay variation      : 26 usec
Best case two-way delay              : 519 usec
Worst case two-way delay             : 650 usec

Loss measurement statistics:
Index      Near-end      Far-end       Near-end      Far-end
```



```
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)        (CIR)         (EIR)        (EIR)
  1          10           8             5            12
  2          12           7             6            16
  3          7            5             0            2
  4          9            6             8            2
  5          3            6             6            4

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)     : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)      : 6.2%
Near-end best case loss (CIR)         : 3
Near-end best case loss ratio (CIR)   : 3%
Near-end worst case loss (CIR)        : 9
Near-end worst case loss ratio (CIR)  : 9%
Far-end best case loss (CIR)          : 5
Far-end best case loss ratio (CIR)    : 5%
Far-end worst case loss (CIR)         : 9
Far-end worst case loss ratio (CIR)   : 9%
Average near-end loss (EIR)           : 4
Average near-end loss ratio (EIR)     : 4%
Average far-end loss (EIR)            : 3.4
Average far-end loss ratio (EIR)      : 3.4%
Near-end best case loss (EIR)         : 0
Near-end best case loss ratio (EIR)   : 0%
Near-end worst case loss (EIR)        : 8
Near-end worst case loss ratio (EIR)  : 8%
Far-end best case loss (EIR)          : 2
Far-end best case loss ratio (EIR)    : 2%
Far-end worst case loss (EIR)         : 6
Far-end worst case loss ratio (EIR)   : 6%
```

show oam ethernet connectivity-fault-management mep-statistics

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 1
  CCMs sent                                    : 6550
  CCMs received out of sequence                : 0
  LBMs sent                                    : 0
  Valid in-order LBRs received                 : 0
  Valid out-of-order LBRs received             : 0
  LBRs received with corrupted data            : 0
  LBRs sent                                    : 0
  LTMs sent                                    : 0
  LTMs received                                : 0
  LTRs sent                                    : 0
  LTRs received                                : 0
  Sequence number of next LTM request          : 0
  1DMs sent                                    : 5
  Valid 1DMs received                          : 0
  Invalid 1DMs received                        : 0
  DMMs sent                                    : 5
  DMRs sent                                    : 0
  Valid DMRs received                          : 5
  Invalid DMRs received                        : 0
  SLM sent                                     : 10
  Valid SLM received                           : 20
```



```
  Invalid SLM received                          : 0
  SLR sent                                      : 20
  Valid SLR received                            : 10
  Invalid SLR received                          : 0

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec    Synthetic Loss measurement
statistics:
     SLM packets sent                           : 100
     SLM packets received                       : 0
     SLR packets sent                           : 100
     SLR packets received                       : 0
     Accumulated SLM statistics:
     Local TXFC1 value                          : 100
     Local RXFC1 value                          : 100
     Last Received SLR frame TXFCftc            : 100
     Last Received SLR frame TXFCbtc            : 100
     SLM Frame Loss:
     Frame Loss (far-end)                       : 0 (0.00 %)
     Frame Loss (near-end)                      : 0 (0.00 %)
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1 remote-mep 101**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
  CCMs sent                                     : 7762
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           : 0
  1DMs sent                                     : 5
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
  DMMs sent                                     : 5
  DMRs sent                                     : 0
  Valid DMRs received                           : 5
  Invalid DMRs received                         : 0
```



```
  SLM sent                                      : 10
  Valid SLM received                            : 20
  Invalid SLM received                          : 0
  SLR sent                                      : 20
  Valid SLR received                            : 10
  Invalid SLR received                          : 0

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
             (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec    Synthetic Loss measurement
statistics:
     SLM packets sent                            : 100
     SLM packets received                        : 0
     SLR packets sent                            : 100
     SLR packets received                        : 0
     Accumulated SLM statistics:
     Local TXFC1 value                           : 100
     Local RXFC1 value                           : 100
     Last Received SLR frame TXFCftc             : 100
     Last Received SLR frame TXFCbtc             : 100
     SLM Frame Loss:
     Frame Loss (far-end)                        : 0 (0.00 %)
     Frame Loss (near-end)                       : 0 (0.00 %)
```

show oam ethernet connectivity-fault-management mep-statistics local-mep remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma1 local-mep 121 remote-mep 101
MEP identifier: 1, MAC address: 3c:61:04:25:fa:95
Remote MEP count: 1
  CCMs sent                                     : 875
  CCMs Received                                 : 300
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           :  12
  1DMs sent                                     : 0
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
```

```
Out of sync 1DMs received                   : 0
DMMs sent                                   : 0
Valid DMMs received                         : 0
Invalid DMMs received                       : 0
DMRs sent                                   : 0
Valid DMRs received                         : 0
Invalid DMRs received                       : 0
LMMs sent                                   : 0
Valid LMMs received                         : 0
Invalid LMMs received                       : 0
LMRs sent                                   : 0
Valid LMRs received                         : 0
Invalid LMRs received                       : 0
SLMs sent                                   : 0
Valid SLMs received                         : 0
Invalid SLMs received                       : 0
SLRs sent                                   : 0
Valid SLRs received                         : 0
Invalid SLRs received                       : 0
Valid AISs generated                        : 0
Valid AISs received                         : 0
```


show oam ethernet connectivity-fault-management path-database

Syntax

show oam ethernet connectivity-fault-management path-database
<*host-mac-address*>
<maintenance-association *ma-name*>
<maintenance-domain *domain-name*>

Release Information

Command introduced in Junos OS Release 8.4.

Description

On M7i and M10i with Enhanced CFEB (CFEB-E), M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management path database information for a host configured with an MEP.

Options

host-mac-address—(Optional) Display connectivity fault management path database information for a specified Ethernet host.

maintenance-association *ma-name*—(Optional) Display connectivity fault management path database information for the specified maintenance association.

maintenance-domain *domain-name*—(Optional) Display connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level

view

List of Sample Output

show oam ethernet connectivity-fault-management path-database on page 2079

Output Fields

Table 153 on page 2078 lists the output fields for the **show oam ethernet connectivity-fault-management path-database** command. Output fields are listed in the approximate order in which they appear.

Table 153: show oam ethernet connectivity-fault-management path-database Output Fields

Field Name	Field Description
Linktrace to	MAC address of the remote MEPs in the path.
Interface	Interface identifier.
Maintenance domain name	Maintenance domain name.
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.

Table 153: show oam ethernet connectivity-fault-management path-database Output Fields (continued)

Field Name	Field Description
Maintenance association name	Maintenance association name.
Local Mep	Local MEP identifier.

Sample Output

show oam ethernet connectivity-fault -management path-database

```
user@host> show oam ethernet connectivity–fault-management path-database
maintenance-domain md1 maintenance-association ma1 00:05:85:79:39:ef
Linktrace to 00:05:85:79:39:ef, Interface : ge-3/0/0
      Maintenance Domain: md1, Level: 7
      Maintenance Association: ma1, Local Mep: 201
```

show oam ethernet connectivity-fault-management policer

Syntax **show oam ethernet connectivity-fault-management policer <maintenance-domain** ***md-name*****> <maintenance-association** ***ma-name*****>**

Release Information Command introduced in Junos OS Release 10.0.

Description On M120, M320, MX Series, T320, and T640 routers displays connectivity-fault-management policer statistics.

Options This command has the following options:

maintenance-domain ***md-name***—Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are displayed for all maintenance associations for all maintenance domains.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association. If this option is not specified, policer statistics are displayed for all maintenance associations for given maintenance domain. This option cannot be specified without specifying maintenance-domain name.

Required Privilege Level view

Related Documentation

- clear oam ethernet connectivity-fault-management policer on page 1468

List of Sample Output

show oam ethernet connectivity-fault-management policer on page 2081
show oam ethernet connectivity-fault-management policer maintenence-domain md-name on page 2081
show oam ethernet connectivity-fault-management policer maintenence-domain md-name maintenance-association ma-name on page 2082

Output Fields Table 154 on page 2080 lists the output fields for the **show oam ethernet connectivity-fault-management policer** command. Output fields are listed in the approximate order in which they appear.

Table 154: show oam ethernet connectivity-fault-management policer Output Fields

Field Name	Field Description
Legend for Policer	Describes the symbols used under the **Scope** and **Type** headings: • G - Global scope • S - Service scope • cc - Continuity check (Type)
Maintenance Domain	Displays the maintenance domain name.

Table 154: show oam ethernet connectivity-fault-management policer Output Fields (continued)

Field Name	Field Description
Level	Displays the maintenance domain level configured.
Maintenance association	Displays the maintenance association name.
Policer	Displays the policer name.
Type	Policer type. Value **cc** means this policer is used only to police continuity check CFM messages. Value **other** means this policer is used only to police non-continuity check CFM messages. Value **all** means this policer is used to police all CFM messages.
Scope	Policer scope. Displays whether the *global* (G) policer configuration is applicable or the session (S) specific policer config is applicable.
Drop count	Displays the number of packets dropped by the indicated policer.

Sample Output

show oam ethernet connectivity-fault -management policer

Displays the policer information for all maintenance associations and their maintenance domains.

```
show oam ethernet connectivity-fault-management policer
Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md1 Level: 1
Maintenance association Policer      Type    Scope Drop count
  ma1                   cfm-policer1 all     G             300
  ma1-2                 cfm-policer1 cc      S             259
  ma1-2                 cfm-policer1 other   G             300
Maintenance Domain: md2 Level: 2
Maintenance association Policer      Type    Scope Drop count
  ma2                   cfm-policer1 cc      G             300
  ma2                   cfm-policer2 other   S             223
```

show oam ethernet connectivity-fault -management policer maintenance-domain

md-name

Displays the policer information for the specified maintenance domain and its maintenance associations.

```
show oam ethernet connectivity-fault-management policer maintenance-domain md1
Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md1 Level: 1
Maintenance association Policer      Type    Scope Drop count
  ma1                   cfm-policer1 all     G            300
  ma1-2                 cfm-policer1 cc      S            259
  ma1-2                 cfm-policer1 other   G            300
```

show oam ethernet connectivity-fault -management policer maintenance-domain md-name maintenance-association ma-name

Displays the policer information for the specified **maintenance-domain** ***md-name*** and **maintenance-association** ***ma-name***.

```
show oam ethernet connectivity-fault-management policer maintenance-domain md5
maintenance-association ma5
Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md5 Level: 5
Maintenance association Policer       Type    Scope Drop count
  ma5                   cfm-policer   cc      S            187
  ma5                   cfm-policer-2 other   S            234
```

show oam ethernet connectivity-fault-management sla-iterator-statistics

Syntax

show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain *md-name*
maintenance-association *ma-name*
sla-iterator *sla-iterator*
<**local-mep** *local-mep-id*>
<**remote-mep** *remote-mep-id*>

Release Information

Command introduced in Junos OS Release 11.4 for EX Series switches.
Command introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 12.2 for ACX Series routers.
Command introduced in Junos OS Release 13.2 for MX Series routers (not on MPC3E Hyperion cards).

Description

Display the Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics.

Options

maintenance-domain *md-name*—Name of an existing connectivity fault management (CFM) maintenance domain.

maintenance-association *ma-name*—Name of an existing CFM maintenance association.

sla-iterator *sla-iterator*— Name of the iterator profile.

local-mep *local-mep-id*—(Optional) Numeric identifier of the local MEP. The range of values is **1** through **8191**.

remote-mep *remote-mep-id*—(Optional) Numeric identifier of the remote MEP. The range of values is **1** through **8192**.

Required Privilege Level

view

Related Documentation

- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- *clear oam ethernet connectivity-fault-management sla-iterator-statistics*

List of Sample Output

show oam ethernet connectivity-fault-management sla-iterator-statistics on page 2086
show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) on page 2086

Output Fields

Table 155 on page 2084 lists the output fields for the **show oam ethernet connectivity-fault-management sla-iterator-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 155: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields

Output Field Name	Output Field Description
Maintenance domain	Name of the maintenance domain.
Level	Level of the maintenance domain level configured.
Maintenance association	Name of the maintenance association.
Local MEP id	Numeric identifier of the local MEP.
Remote MEP id	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
Delay weight	Calculation weight of delay.
Delay variation weight	Calculation weight of delay variation.
DMM sent	Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.
DMM skipped for threshold hit	Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.
DMM skipped for threshold hit window	Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
DMR received	Number of delay measurement reply (DMR) frames received.
DMR out of sequence	Total number of DMR out of sequence packets received.
DMR received with invalid time stamps	Total number of DMR frames received with invalid timestamps.

Table 155: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (continued)

Output Field Name	Output Field Description
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Average one-way forward delay variation	Average one-way forward delay variation for the statistics displayed in microseconds.
Average one-way backward delay variation	Average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average two-way delay	Weighted average two-way delay for the statistics displayed in microseconds.
Weighted average two-way delay variation	Weighted average two-way delay variation for the statistics displayed in microseconds.
Weighted average one-way backward delay variation	Weighted average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average one-way forward delay variation	Weighted average one-way forward delay variation for the statistics displayed in microseconds.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLM packets received	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
SLR packets sent	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFCl at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t	Value of the local counter RxFCl at the time of SLR frame transmission.

Table 155: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (continued)

Output Field Name	Output Field Description
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault -management sla-iterator-statistics

```
user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics
sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1
remote-mep 2
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent                                         : 23898520
DMM skipped for threshold hit                    : 11000
DMM skipped for threshold hit window             : 0
DMR received                                     : 23851165
DMR out of sequence                              : 1142
DMR received with invalid time stamps            : 36540
Average two-way delay                            : 129 usec
Average two-way delay variation                  : 15 usec
Average one-way forward delay variation          : 22 usec
Average one-way backward delay variation         : 22 usec
Weighted average two-way delay                   : 134 usec
Weighted average two-way delay variation         : 8 usec
Weighted average one-way forward delay variation  : 6 usec
Weighted average one-way backward delay variation : 2 usec
```

Sample Output

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers)

```
user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain mdu maintenance-association mau local-mep 4 remote-mep 3 sla-iterator
lm
Iterator statistics:
  Maintenance domain: 2, Level: 2
  Maintenance association: W-160432000-001, Local MEP id: 2
  Remote MEP id: 1, Remote MAC address: 00:90:69:0a:43:39
  Iterator name: iter1, Iterator Id: 1
  Iterator cycle time: 100ms, Iteration period: 10 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2012-09-25 02:15:31 PDT (00:00:45 ago)
```



```
Reset reason: Adjacency flap
Iterator loss measurement statistics:
  LMM sent                          : 444
  LMM skipped for threshold hit     : 0
  LMM skipped for threshold hit window: 0
  LMR received                      : 444
  LMR out of sequence               : 0
  LMR forwarding-class mismatch     : 0
Accumulated transmit statistics:
  Near-end (CIR)                    : 0
  Far-end (CIR)                     : 0
  Near-end (EIR)                    : 0
  Far-end (EIR)                     : 0
Accumulated receive statistics:
  Near-end (CIR)                    : 0
  Far-end (CIR)                     : 0
  Near-end (EIR)                    : 0
  Far-end (EIR)                     : 0
Accumulated loss statistics:
  Near-end loss (CIR)               : 0
  Near-end loss-ratio (CIR)         : 0 (0.00000%)
  Far-end loss (CIR)                : 0
  Far-end loss-ratio (CIR)          : 0 (0.00000%)
  Near-end loss (CIR)               : 0
  Near-end loss-ratio (EIR)         : 0 (0.00000%)
  Far-end loss (CIR)                : 0
  Far-end loss-ratio (EIR)          : 0 (0.00000%)
Last loss measurement statistics:
  Near-end (CIR)                    : 0
  Far-end (CIR)                     : 0
  Near-end (EIR)                    : 0
  Far-end (EIR)
```


show oam ethernet connectivity-fault-management synthetic-loss-statistics

Syntax

show oam ethernet connectivity-fault-management synthetic-loss-statistics
<local-mep *local-mep-id*>
maintenance-association *ma-name*
<count *entry-count*>
maintenance-domain *md-name*
<remote-mep *remote-mep-id*>

Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.

Description

On MX Series routers with Modular Port Concentrators (MPCs) with Ethernet interfaces, display the on-demand ETH-SLM statistics.

Options

count *entry-count*—(Optional) Number of entries to display from the statistics table. The range of values is from 1 through 100. The default value is 100.

local-mep *local-mep-id*—(Optional) Numeric identifier of the local MEP. The range of values is from **1** through **8192**.

maintenance-association *ma-name*—Name of an existing CFM maintenance association.

maintenance-domain *md-name*—Name of an existing connectivity fault management (CFM) maintenance domain.

remote-mep *remote-mep-id*—(Optional) Numeric identifier of the remote MEP. The range of values is from **1** through **8192**.

Required Privilege Level

view

Related Documentation

- clear oam ethernet connectivity-fault-management statistics on page 1469
- clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 1471
- show oam ethernet connectivity-fault-management interfaces on page 2037
- show oam ethernet connectivity-fault-management mep-database on page 2055
- show oam ethernet connectivity-fault-management mep-statistics on page 2066

List of Sample Output

show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 2089

Output Fields

Table 156 on page 2089 lists the output fields for the **show oam ethernet connectivity-fault-management synthetic-loss-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 156: show oam ethernet connectivity-fault-management synthetic-loss-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLM packets received	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
SLR packets sent	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TxFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFCl at the time of SLM frame transmission.
Last Received SLR frame TXFCb(t	Value of the local counter RxFCl at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault-management synthetic-loss-statistics

`user@switch>` **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md6 maintenance-association ma6**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
  Synthetic Loss measurement statistics:
   SLM packets sent                            : 100
   SLM packets received                        : 0
   SLR packets sent                            : 100
   SLR packets received                        : 0
  Accumulated SLM statistics:
   Local TXFC1 value                           : 100
   Local RXFC1 value                           : 100
   Last Received SLR frame TXFCftc             : 100
   Last Received SLR frame TXFCbtc             : 100
  SLM Frame Loss:
   Frame Loss (far-end)                        : 0 (0.00 %)
   Frame Loss (near-end)                       : 0 (0.00 %)
```


show oam ethernet evc

Syntax **show oam ethernet evc <*evc-id*>**

Release Information Command introduced in Junos OS Release 9.5.

Description On MX Series routers with OAM Ethernet Virtual Connection (EVC) configurations, displays the EVC configuration and status information.

Options This command has no options.

Required Privilege Level View

Output Fields Table 157 on page 2091 lists the output fields for the **show oam ethernet evc** command. Output fields are listed in the approximate order in which they appear.

Table 157: show oam ethernet evc Output Fields

Field Name	Field Description
EVC identifier	Header for the EVC information showing the EVC name, configuration, and active/inactive status.
UNI count	Number of configured and active UNIs.
Protocol	Protocol configured between the UNIs.
Local UNIs	Heading for the list of local UNIs
UNI Identifier	Name of the UNI.
Interface	Interface type-dpc/pic/port.unit-number.
Status	Status operational or not operational.

Sample Output

show oam ethernet evc

```
user@host> show oam ethernet evc
EVC identifier: evc1, Point-to-Point, Active
  UNI count: Configured(2), Active(2)
  Protocol: cfm, Management domain: md, Management association: ma
  Local UNIs:
    UNI Identifier      Interface     Status
    uni1                ge-1/1/1      Operational
    uni2                ge-1/1/1      Not Operational
```


show oam ethernet fnp interface

Syntax

show oam ethernet fnp interface
<*ethernet-interface-name*>
<routing-instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 11.4.

Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) information for Ethernet interfaces.

Options

interface-name—(Optional) Display Ethernet FNP information for the specified Ethernet interface only.

routing-instance-name—(Optional) Display FNP for the specified routing instance.

Required Privilege Level

view

List of Sample Output

show oam ethernet fnp interface on page 2092

Output Fields

Table 158 on page 2092 lists the output fields for the **show oam ethernet fnp interface** command. Output fields are listed in the approximate order in which they appear.

Table 158: show oam ethernet fnp interface Output Fields

Field Name	Field Description
Interface	Name of the interface for the displayed information.
VLAN	Name of the VLAN.
State	Displays state of the interface.
FNP Message Interface	Displays the message interface type.
FNP Message Source MAC	Displays the source MAC address.

Sample Output

show oam ethernet fnp interface

```
user@host> show oam ethernet fnp interface
  The FNP controlled interfaces are:
  Interface    VLAN   State   FNP message    FNP message
                              Interface      Source MAC

  ge-0/0/0.30   30     down   lsi.1054976    a0:aa:aa:aa:aa:aa
```

```
ge-0/0/0.20  20     down   lsi.1054976   a0:aa:aa:aa:aa:aa
```

show oam ethernet fnp messages

Syntax

show oam ethernet fnp messages
<interface *interface-name*>
<routing instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 11.4

Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) messages.

Options

interface-name—(Optional) Display Ethernet FNP messages for the specified Ethernet interface only.

routing-instance-name—(Optional) Display FNP messages for the specified routing instance.

Required Privilege Level

view

List of Sample Output

show oam ethernet fnp messages on page 2095

Output Fields

Table 159 on page 2094 lists the output fields for the **show oam ethernet fnp messages** command. Output fields are listed in the approximate order in which they appear.

Table 159: show oam ethernet fnp messages Output Fields

Field Name	Field Description
Message from source MAC address	The source MAC address of the message.
Originating port number	Port number of the original message.
Time since last message	Elapsed time in hours, minutes, and seconds since the last message was received.
Time since last message update	Elapsed time in hours, minutes, and seconds since the last message was updated.
Total messages received	Number of messages received.
Domain ID	Domain ID of the message.
STP Root ID	The spanning tree Root ID of the message.
Trigger Reason	The reason why the message was triggered.
Effected VLANs	Number of VLANs that are affected.

Table 159: show oam ethernet fnp messages Output Fields (continued)

Field Name	Field Description
Disabled interfaces	Name of the interfaces that are disabled.

Sample Output

show oam ethernet fnp messages

```
user@host> show oam ethernet fnp messages
  Active FNP messages on interface lsi.1054465
  Message source MAC: a0:aa:aa:aa:aa:aa
    Originating port number: 141077
    Time since last message: 00:00:00
    Time since last message update: 00:00:00
    Total messages received: 1
    Domain ID: 0
    STP Root ID: 0.f0:ff:ff:ff:ff:ff
    Trigger reason: todo
    Effected VLANs: 10
    Disabled interfaces:
      Interface VLAN
      ge-0/0/0.10 10
```


show oam ethernet fnp status

Syntax

show oam ethernet fnp status
<interface *interface-name*>
<routing instance *routing-instance-name*>

Release Information

Command introduced in Junos OS Release 11.4

Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) status.

Options

interface-name—(Optional) Display Ethernet FNP information for the specified Ethernet interface only.

routing-instance-name—(Optional) Display FNP for the specified routing instance.

Required Privilege Level

view

List of Sample Output

show oam ethernet fnp status on page 2096

Output Fields

Table 160 on page 2096 lists the output fields for the **show oam ethernet fnp status** command. Output fields are listed in the approximate order in which they appear.

Table 160: show oam ethernet fnp status Output Fields

Field Name	Field Description
FNP interval	The time interval between messages.
Loss threshold	The number of messages that can be lost before FNP is marked as down.
FNP enabled interfaces	Displays interfaces that are enabled.
Interface	The name of the interface.
Domain ID	Domain ID of the message.
STP Root ID	The spanning tree Root ID of the message.
FNP Messages	The total number of messages received.

Sample Output

show oam ethernet fnp status

`user@host>` **show oam ethernet status**

```
FNP interval:
Loss threshold
FNP enabled interfaces
Interface      Domain ID      STP Root ID             FNP Messages
ge-0/0/0.1278        100      0.f0:ff:ff:ff:ff:ff                0
```

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 2103
show oam ethernet link-fault-management brief (Loopback tracking) on page 2103
show oam ethernet link-fault-management detail on page 2103
show oam ethernet link-fault-management detail (backup Routing Engine) on page 2104

Output Fields

Table 161 on page 2098 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 161: 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

Table 161: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
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
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

Table 161: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
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
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**

Table 161: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
Errors in period	The number of symbol errors in the period reported in the received event PDU.	**detail**
Total errors	The number of errored symbols that have been reported in received event TLVs 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**

Table 161: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
Total errors	The number of errored symbols reported in event TLVs that have been transmitted since the OAM sublayer was reset.	**detail**
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 oam ethernet link-fault-management detail (backup Routing Engine)

```
user@host> show oam ethernet link-fault-management ge-0/2/0 detail
Interface: ge-0/2/0
    Status: Running, Discovery state: Send Any
    Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms
    Peer address: ac:4b:c8:81:90:a4
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM flags receive statistics:
      Critical event: 0, Dying gasp: 0, Link fault: 0
    OAM transmit statistics:
      Information: 0, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 786, 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 received frame seconds 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
    Loopback tracking: Enabled, Loop status: Not Found
    Detect LOC: Enabled, LOC status: Not Found
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: unsupported, Link events: supported
      Variable requests: unsupported
Application profile statistics:
  Profile Name                  Invoked      Executed
  LK_ADJ_LOSS100_1                        1         1
  LK_ADJ_LOSS100_2                        1         0
  LK_ADJ_LOSS100_3                        1         0
  LK_ADJ_LOSS101_1                        1         1
  LK_ADJ_LOSS101_2                        1         0
  LK_ADJ_LOSS101_3                        1         0
```



```
LK_ADJ_LOSS106_1                               0          0
LK_ADJ_LOSS106_2                               0          0
LK_ADJ_LOSS106_3                               0          0
LK_ADJ_LOSS107_1                               0          0
LK_ADJ_LOSS107_2                               0          0
LK_ADJ_LOSS107_3                               0          0
```


show oam ethernet lmi

Syntax **show oam ethernet lmi (interface <*interface-name*>)**

Release Information Command introduced in Junos OS Release 9.5.

Description On routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet, and OAM Ethernet Local Management Interface (E-LMI) configuration, display the LMI information for the configured interfaces or optionally for a specified interface.

NOTE: On MX Series routers, E-LMI is supported on Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), and Aggregated Ethernet (ae) interfaces configured on MX Series routers with DPC only.

Options **interface**—(Optional) Display LMI information for a specified interface.

interface-name—(Optional) Display Ethernet LMI information for the specified interface only.

Required Privilege Level View

Output Fields Table 162 on page 2106 lists the output fields for the **show oam ethernet lmi** command. Output fields are listed in the approximate order in which they appear.

Table 162: show oam ethernet lmi Output Fields

Field Name	Field Description
Physical Interface	Header for the EVC information showing the Ethernet virtual circuit (EVC) name, configuration, and active/inactive status.
UNI Identifier	Name of the UNI.
EVC map type	EVC configuration.
Polling verification timer	Polling verification timer status.
E-LMI state	Operational status of the E-LMI configuration in the interfaces or specified interface.
Priority/Untagged VLAN ID	To be provided.
Default EVC	The EVC set as the default EVC.
Associated EVCs	Heading for the list of configured EVCs.

Table 162: show oam ethernet lmi Output Fields (continued)

Field Name	Field Description
EVC Indentifier	EVC name.
Reference ID	To be provided.
Status	Status active or not active.
CE VLAN IDs	Customer edge VLAN ID numbers.

Sample Output

show oam ethernet lmi interface

```
user@host> show oam ethernet lmi interface ge-1/1/1
Physical interface: ge-1/1/1, Physical link is Up
  UNI identifier: uni-ce1, EVC map type: Bundling
  Polling verification timer: Enabled, E-LMI state: Operational
  Priority/Untagged VLAN ID: 20, Default EVC: evc1
  Associated EVCs:
    EVC         Reference     Status           CE VLAN IDs
    Identifier  ID
    evc1            1        Active (New)        1-2048
    evc2             2        Not Active          2049-4096
```

show oam ethernet lmi statistics

Syntax **show oam ethernet lmi statistics <interface *interface-name*>**

Release Information Command introduced in Junos OS Release 9.5.

Description On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Local Management Interface (LMI) statistics.

Options **interface**—(Optional) Display LMI statistics for a specified interface.

interface-name—(Optional) Display Ethernet LMI information for the specified Ethernet interface only.

Required Privilege Level view

List of Sample Output show oam ethernet lmi statistics on page 2109

Output Fields Table 163 on page 2108 lists the output fields for the **show oam ethernet lmi statistics** command. Output fields are listed in the approximate order in which they appear.

Table 163: show oam ethernet lmi statistics Output Fields

Field Name	Field Description
Physical interface	Name of the interface for the displayed statistics.
Reliability errors	Number of E-LMI reliability errors logged.
Protocol errors	Number of E-LMI protocol errorse.
Status check received	Number of E-LMI status check receive errors.
Status check sent	Number of E-LMI status check sent errors.
Full status received	Number of E-LMI full status receive errors.
Full status sent	Number of E-LMI full status sent errors.
Full status continued received	Number of E-LMI status continued received errors.
Full status continued sent	Number of E-LMI full status continued sent errors.
Asynchronous status sent	Number of E-LMI asynchronous status sent errors.

Sample Output

show oam ethernet lmi statistics

```
user@host> show oam ethernet lmi statistics interface ge-1/1/1
Physical interface: ge-1/1/1
  Reliability errors                      4  Protocol errors
 0
  Status check received                   0  Status check sent
 0
  Full status received                  694  Full status sent
694
  Full status continued received          0  Full status continued sent
 0
  Asynchronous status sent                0
```

show pppoe interfaces

Syntax

show pppoe interfaces
<brief | detail
<pp0.*logical*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display session-specific information about PPPoE interfaces.

Options

none—Display interface information for all PPPoE interfaces.

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

pp0.*logical*—(Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16385. The logical unit number for dynamic interfaces can be a value from 1073741824 through the maximum number of logical interfaces supported on your router.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show pppoe interfaces on page 2112
show pppoe interfaces (Status for the Specified Interface) on page 2112
show pppoe interfaces brief on page 2112
show pppoe interfaces detail on page 2113
show pppoe interfaces (PPPoE Subscriber Interface with ACI Interface Set) on page 2113

Output Fields

Table 164 on page 2110 lists the output fields for the **show pppoe interfaces** command. Output fields are listed in the approximate order in which they appear. Not all fields are displayed for PPPoE interfaces on M120 and M320 routers in server mode.

Table 164: show pppoe interfaces Output Fields

Field Name	Field Description	Level of Output
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
State	State of the logical interface: **up** or **down**.	All levels
Session ID	Session ID.	All levels

Table 164: show pppoe interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Type	Origin of the logical interface: **Static** or **Dynamic**. Indicates whether the interface was statically or dynamically created.	**detail extensive** none
Service name	Type of service required (can be used to indicate an ISP name or a class or quality of service).	**detail extensive** none
Configured AC name	Configured access concentrator name.	**detail extensive** none
Session AC name	Name of the access concentrator.	**detail extensive** none
Remote MAC address or **Remote MAC**	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.	All levels
Session uptime	Length of time the session has been up, in *hh:mm:ss*.	**detail extensive** none
Dynamic Profile	Name of the dynamic profile that was used to create this interface. If the interface was statically created, this field is not displayed.	**detail extensive** none
Underlying interface	Interface on which PPPoE is running.	All levels
Agent Circuit ID	Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. If the agent circuit ID is not configured, this field is not displayed.	**detail extensive** none
Agent Remote ID	Agent remote identifier that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. If the agent remote ID is not configured, this field is not displayed.	**detail extensive** none
ACI Interface Set	Internally-generated name of the dynamic ACI interface set, if configured, and the set index number of the ACI entry in the session database.	**detail extensive** none

Table 164: show pppoe interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Packet Type	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.	**extensive**

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 66
  State: Down, Session ID: None,
  Service name: None, Configured AC name: sapphire,
  Session AC name: None, Remote MAC address: 00:00:5e:00:53:00,
  Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
  Underlying interface: at-5/0/0.0 Index 71
```

show pppoe interfaces (Status for the Specified Interface)

```
user@host> show pppoe interfaces pp0.1073741827
pp0.1073741827 Index 70
  State: Session Up, Session ID: 30, Type: Dynamic,
  Session AC name: velorum,
  Remote MAC address: 00:00:5e:00:53:c1,
  Session uptime: 16:45:46 ago,
  Underlying interface: ge-2/0/3.1 Index 73
  Service name: premium
  Dynamic Profile: PppoeProfile
  Agent Circuit ID: velorum-ge-2/0/3
  Agent Remote ID: westford
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
Interface       Underlying            State       Session    Remote
                interface                         ID         MAC
pp0.0           ge-2/0/3.2            Session Up  27         00:00:5e:00:53:c1
pp0.1           ge-2/0/3.2            Session Up  28         00:00:5e:00:53:c1
pp0.1073741824  ge-2/0/3.1            Session Up  29         00:00:5e:00:53:c1
```

```
pp0.1073741825  ge-2/0/3.1           Session Up  30          00:00:5e:00:53:c1
pp0.1073741826  ge-2/0/3.1           Session Up  31          00:00:5e:00:53:c1
```

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 66
  State: Down, Session ID: None, Type: Static,
  Service name: None, Configured AC name: sapphire,
  Session AC name: None, Remote MAC address: 00:00:5e:00:53:00,
  Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
  Underlying interface: at-5/0/0.0 Index 71
```

show pppoe interfaces (PPPoE Subscriber Interface with ACI Interface Set)

```
user@host> show pppoe interfaces pp0.1073741827
pp0.1073741827 Index 346
  State: Session Up, Session ID: 4, Type: Dynamic,
  Service name: AGILENT, Remote MAC address: 00:00:5e:00:53:62,
  Session AC name: nbc,
  Session uptime: 6d 02:22 ago,
  Dynamic Profile: aci-vlan-pppoe-profile,
  Underlying interface: demux0.1073741826 Index 345
  Agent Circuit ID: aci-ppp-dhcp-dvlan-50
  ACI Interface Set: aci-1002-demux0.1073741826 Index 2
```

show pppoe service-name-tables

Syntax

show pppoe service-name-tables
<*table-name*>

Release Information

Command introduced in Junos OS Release 10.0.

Description

Display configuration information about PPPoE service name tables.

Options

none—Display the names of configured PPPoE service name tables.

table-name—(Optional) Name of a configured PPPoE service name table.

Required Privilege Level

view

Related Documentation

- Verifying a PPPoE Configuration on page 360
- *Verifying and Managing Dynamic PPPoE Configuration*

List of Sample Output

show pppoe service-name-tables on page 2115
show pppoe service-name-tables (For the Specified Table Name) on page 2115

Output Fields

Table 165 on page 2114 lists the output fields for the **show pppoe service-name-tables** command. Output fields are listed in the approximate order in which they appear.

Table 165: show pppoe service-name-tables Output Fields

Field Name	Field Description	Level of Output
Service Name Table	Name of the PPPoE service name table.	none
Service Name	Name of a configured service in the PPPoE service name table: • **<empty>**—Service of zero length that represents an unspecified service • **<any>**—Default service for non-empty service entries that do not match the configured empty or named service entries • ***service-name***—Named service entry	none
Action	Action taken when the PPPoE underlying interface interface receives a PPPoE Active Discovery Initiation (PADI) packet with the specified named service, **empty** service, **any** service, or ACI/ARI pair: • **Delay *seconds***—Number of seconds that the interface delays before responding with a PPPoE Active Discovery Offer (PADO) packet • **Drop**—Interface drops (ignores) the packet. • **Terminate**—Interface responds immediately with a PADO packet	none

Table 165: show pppoe service-name-tables Output Fields (continued)

Field Name	Field Description	Level of Output
Dynamic Profile	Name of the dynamic profile with which the router creates a dynamic PPPoE subscriber interface. A dynamic profile can be assigned to a named service, **empty** service, **any** service, or ACI/ARI pair.	none
Routing Instance	Name of the routing instance in which to instantiate the dynamic PPPoE subscriber interface. A routing instance can be assigned to a named service, **empty** service, **any** service, or ACI/ARI pair.	none
Max Sessions	Maximum number of active PPPoE sessions that the router can establish with the specified named service, **empty** service, or **any** service.	none
Active Sessions	Current count of active PPPoE sessions created using the specified named service, **empty** service, or **any** service. The Active Sessions value cannot exceed the Max Sessions value.	none
ACI	Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. An ACI can be configured as part of an ACI/ARI pair for a named service, **empty** service, or **any** service.	none
ARI	Agent remote identifier (ARI) that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. An ARI can be configured as part of an ACI/ARI pair for a named service, **empty** service, or **any** service.	none
Static Interface	Name of the static PPPoE interface reserved for exclusive use by the PPPoE client with matching ACI/ARI information. A static interface can be configured only for an ACI/ARI pair.	none

Sample Output

show pppoe service-name-tables

```
user@host> show pppoe service-name-tables
Service Name Table: test1
Service Name Table: test2
Service Name Table: test3
```

show pppoe service-name-tables (For the Specified Table Name)

```
user@host> show pppoe service-name-tables Table1
Service Name Table: Table1
  Service Name: <empty>
    Action: Terminate
    Dynamic Profile: BasicPppoeProfile
    Max Sessions: 100
    Active Sessions: 3
  Service Name: <any>
    Action: Drop
    ACI: velorum-ge-2/0/3
```



```
  ARI: westford
    Action: Terminate
    Static Interface: pp0.100
  ACI: volantis-ge-5/0/5
  ARI: sunnyvale
     Action: Terminate
     Static Interface: pp0.101
Service Name: Wholesale
  Action: Terminate
  Dynamic Profile: WholesalePppoeProfile
  Routing Instance: WholesaleRI
  Max Sessions: 16000
  Active Sessions: 4
```


show pppoe sessions

Syntax

show pppoe sessions
<aci *circuit-id-string*>
<ari *remote-id-string*>
<service *service-name*>

Release Information

Command introduced in Junos OS Release 10.2.

Description

Display information about all active PPPoE sessions on the router, or about the active PPPoE sessions established for a specified service name, agent circuit identifier (ACI), or agent remote identifier (ARI).

Options

none—Display information for all active PPPoE sessions on the router.

aci *circuit-id-string*—(Optional) Display information only for active PPPoE sessions established with the specified agent circuit identifier. The agent circuit identifier corresponds to the DSLAM interface that initiated the service request.

ari *remote-id-string*—(Optional) Display information only for active PPPoE sessions established with the specified agent remote identifier. The agent remote identifier corresponds to the subscriber associated with the DSLAM interface that initiated the service request.

service *service-name*—(Optional) Display information only for active PPPoE sessions established with the specified service, where *service-name* can be **empty**, **any**, or a named service.

Required Privilege Level

view

Related Documentation

- Verifying a PPPoE Configuration on page 360
- *Verifying and Managing Dynamic PPPoE Configuration*

List of Sample Output

show pppoe sessions (For All Active Sessions) on page 2118
show pppoe sessions (For All Active Sessions Matching the Agent Circuit Identifier) on page 2118

Output Fields

Table 166 on page 2117 lists the output fields for the **show pppoe sessions** command. Output fields are listed in the approximate order in which they appear.

Table 166: show pppoe sessions Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the statically-created or dynamically-created PPPoE interface for the active PPPoE session.	none

Table 166: show pppoe sessions Output Fields (continued)

Field Name	Field Description	Level of Output
Underlying interface	Interface on which PPPoE is running.	none
State	State of the PPPoE session; displays **Session Up** for active PPPoE sessions.	none
Session ID	PPPoE session identifier.	none
Remote MAC	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.	none

Sample Output

show pppoe sessions (For All Active Sessions)

```
user@host> show pppoe sessions
Interface         Underlying            State         Session    Remote
                  interface                           ID         MAC
pp0.0             ge-2/0/3.2            Session Up    27         00:00:5e:00:53:c1
pp0.1             ge-2/0/3.2            Session Up    28         00:00:5e:00:53:c1
pp0.1073741824    ge-2/0/3.1            Session Up    29         00:00:5e:00:53:c1
pp0.1073741825    ge-2/0/3.1            Session Up    30         00:00:5e:00:53:c1
pp0.1073741826    ge-2/0/3.1            Session Up    31         00:00:5e:00:53:c1
```

show pppoe sessions (For All Active Sessions Matching the Agent Circuit Identifier)

```
user@host> show pppoe sessions aci “velorum-ge-2/0/3”
Interface         Underlying            State         Session    Remote
                  interface                           ID         MAC
pp0.0             ge-2/0/3.2            Session Up    27         00:00:5e:00:53:c1
pp0.1             ge-2/0/3.2            Session Up    28         00:00:5e:00:53:c1
```


show pppoe statistics

Syntax

show pppoe statistics
<*logical-interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
logical-interface-name option introduced in Junos OS Release 10.1.

Description

Display statistics information about PPPoE interfaces.

Options

none—Display PPPoE statistics for all interfaces.

logical-interface-name—(Optional) Name of a PPPoE underlying logical interface.

Required Privilege Level

view

Related Documentation

- *show ppp address-pool*
- show pppoe underlying-interfaces on page 2121

List of Sample Output

show pppoe statistics on page 2120
show pppoe statistics (For the Specified Underlying Interface Only) on page 2120

Output Fields

Table 167 on page 2119 lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 167: show pppoe statistics Output Fields

Field Name	Field Description
Active PPPoE sessions	Total number of active PPPoE sessions and the number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.

Table 167: show pppoe statistics Output Fields (continued)

Field Name	Field Description
Timeouts	Information about timeouts that occurred during the PPPoE session (not displayed for M120, M320, and MX Series routers): • **PADI**—No PADR packet has been received within the timeout period. (This value is always zero and is not supported.) • **PADO**—No PPPoE Active Discovery Offer packet has been received within the timeout period. • **PADR**—No PADS packet has been received within the timeout period.

Sample Output

show pppoe statistics

```
user@host> show pppoe statistics
Active PPPoE sessions: 1
  PacketType          Sent      Received
  PADI                   0             0
  PADO                   0             0
  PADR                   0             0
  PADS                   0             0
  PADT                   0             0
  Service name error     0             0
  AC system error        0             0
  Generic error          0             0
  Malformed packets      0             0
  Unknown packets        0             0
  Timeouts
  PADI                   0
  PADO                   0
  PADR                   0
```

show pppoe statistics (For the Specified Underlying Interface Only)

```
user@host> show pppoe statistics ge-4/0/3.2
Active PPPoE sessions: 4
  PacketType                   Sent         Received
    PADI                          0                5
    PADO                          5                0
    PADR                          0                5
    PADS                          4                0
    PADT                          0                1
    Service name error            0                0
    AC system error               0                0
    Generic error                 0                0
    Malformed packets             0                0
    Unknown packets               0                0
```


show pppoe underlying-interfaces

Syntax

show pppoe underlying-interfaces
<brief | detail | extensive>
<lockout>
<*logical-interface-name*>

Release Information

Command introduced in Junos OS Release 10.0.
lockout option added in Junos OS Release 11.4.

Description

Display information about PPPoE underlying interfaces.

Options

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

lockout—(Optional) Display summary information about the lockout condition and the lockout grace period for PPPoE clients on the PPPoE underlying interface.

logical-interface-name—(Optional) Name of a PPPoE underlying logical interface.

Required Privilege Level

view

Related Documentation

- *Verifying and Managing Dynamic PPPoE Configuration*
- *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*
- *Configuring the PPPoE Family for an Underlying Interface*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

List of Sample Output

Output Fields

Table 168 on page 2122 lists the output fields for the **show pppoe underlying-interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 168: show pppoe underlying-interfaces Output Fields

Field Name	Field Description	Level of Output
Underlying Interface	Name of the PPPoE underlying logical interface.	All levels
Service Name Table	Name of the service name table.	All levels
Dynamic Profile	Name of the dynamic profile that was used to create this interface. If the interface was statically created, then the value is **none**.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive**
State	Origin of the logical interface: **Static** or **Dynamic**. Indicates whether the interface was statically or dynamically created.	**detail extensive**
Operational States	Fields in this block are actual operational values rather than simply the configured values. The operational values can be the result of RADIUS-initiated changes.	**detail extensive**
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface. When this number of logical interfaces has been established, all subsequent PPPoE Active Discovery Initiation (PADI) packets are dropped and all subsequent PPPoE Active Discovery Request (PADR) packets trigger PPPoE Active Discovery Session (PADS) error responses.	**detail extensive**
Max Sessions VSA Ignore	Whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configure with the **max-sessions** statement: **Off** (default) or **On**.	**detail extensive** none
Active Sessions	Number of active PPPoE sessions on the underlying interface. If a dynamic profile is listed, then it is the number of active PPPoE sessions on the underlying interface that are using this profile. The Active Sessions value must not exceed the Max Sessions value.	**detail extensive**
Agent Circuit Identifier	Whether the underlying interface is configured with the **agent-circuit-identifier** statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on agent circuit identifier (ACI) information. **Autosensing** indicates that creation of ACI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ACI-based dynamic VLANs is not configured on the underlying interface, this field does not appear. NOTE: The Agent Circuit Identifier field is replaced with the Line Identity field when an ALI interface set is configured with the **line-identity** autoconfiguration stanza.	**detail extensive** none

Table 168: show pppoe underlying-interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Line Identity	Whether the underlying interface is configured with the **line-identity** statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on the specified trusted option: ACI, ARI, both, or neither. **Autosensing** indicates that creation of ALI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ALI dynamic VLANs based on trusted options is not configured on the underlying interface, this field does not appear. NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the **agent-circuit-id** autoconfiguration stanza.	**detail extensive** none
Duplicate Protection	State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client. The uniqueness of the PPPoE client is determined by the client's MAC address.	**detail extensive**
Short Cycle Protection	State of PPPoE short cycle protection: **mac-address**, **circuit-id**, or **Off**. Enabling short cycle protection, also known as PPPoE lockout, on the PPPoE underlying interface temporarily prevents (locks out) a failed or short-lived (short-cycle) PPPoE subscriber session from reconnecting to the router for a default or configurable period of time. PPPoE client sessions are identified by their unique media access control (MAC) source address or agent circuit identifier (ACI) value.	**detail extensive**
Direct Connect	State of the configuration to ignore DSL Forum VSAs: **On** or **Off**. When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	**detail extensive** none
AC Name	Name of the access concentrator.	**detail extensive**
PacketType	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • **PADI**—PPPoE Active Discovery Initiation packets. • **PADO**—PPPoE Active Discovery Offer packets. • **PADR**—PPPoE Active Discovery Request packets. • **PADS**—PPPoE Active Discovery Session-Confirmation packets. • **PADT**—PPPoE Active Discovery Termination packets. • **Service name error**—Packets for which the Service-Name request could not be honored. • **AC system error**—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • **Generic error**—Packets that indicate an unrecoverable error occurred. • **Malformed packets**—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • **Unknown packets**—Unrecognized packets.	**detail extensive**

Table 168: show pppoe underlying-interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Lockout Time (sec)	The PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period if **Short Cycle Protection** is enabled (**On**): • **Min**—Minimum lockout time, in seconds, configured on the PPPoE underlying interface. • **Max**—Maximum lockout time, in seconds, configured on the PPPoE underlying interface. • **Total clients in lockout**—Number of PPPoE clients currently undergoing lockout. • **Total clients in lockout grace period**—Number of PPPoE clients currently in a lockout grace period. A *lockout grace period* occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.	**extensive**
Client Address	MAC source address of the PPPoE client.	**extensive**
Current	Current lockout time, in seconds; displays **0** (zero) if the PPPoE client is not undergoing lockout.	**extensive**
Elapsed	Time elapsed into the lockout period, in seconds; displays **0** if the PPPoE client is not undergoing lockout	**extensive**
Next	Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.	**extensive**

Sample Output

show pppoe underlying-interfaces brief

```
user@host> show pppoe underlying-interfaces brief
Underlying Interface     Service Name Table     Dynamic Profile
ge-4/0/3.1               Premium                None
ge-4/0/3.2               None                   PppoeProfile
```

show pppoe underlying-interfaces detail

```
user@host> show pppoe underlying-interfaces detail
ge-4/0/3.1 Index 73
  Operational States:
  State: Static, Dynamic Profile: None,
  Max Sessions: 4000, Max Sessions VSA Ignore: Off,
  Active Sessions: 0,
  Service Name Table: Premium,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: On,
  Short Cycle Protection: Off

ge-4/0/3.2 Index 78
  Operational States:
  State: Dynamic, Dynamic Profile: PppoeProfile,
  Max Sessions: 500, Max Sessions VSA Ignore: Off,
  Active Sessions: 3,
```

```
  Service Name Table: None,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: On,
  Short Cycle Protection: Off
```

show pppoe underlying-interfaces extensive

```
user@host> show pppoe underlying-interfaces extensive
ge-4/0/3.1 Index 73
  Operational States:
  State: Static, Dynamic Profile: None,
  Max Sessions: 4000, Max Sessions VSA Ignore Off,
  Active Sessions: 0,
  Service Name Table: None,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: Off,
  Short Cycle Protection: Off

  PacketType                       Sent         Received

    PADI                              0                0
    PADO                              0                0
    PADR                              0                0
    PADS                              0                0
    PADT                              0                0
    Service name error                0                0
    AC system error                   0                0
    Generic error                     0                0
    Malformed packets                 0                0
    Unknown packets                   0                0

ge-4/0/3.2 Index 78
  Operational States:
  State: Dynamic, Dynamic Profile: PppoeProfile,
  Max Sessions: 4000, Max Sessions VSA Ignore: Off
  Active Sessions: 3,
  Service Name Table: None,
  Direct Connect: Off,
  AC Name: velorum, Duplicate Protection: Off,
  Short Cycle Protection: Off

  PacketType                       Sent         Received
    PADI                              0                5
    PADO                              5                0
    PADR                              0                5
    PADS                              4                0
    PADT                              0                1
    Service name error                0                0
    AC system error                   0                0
    Generic error                     0                0
    Malformed packets                 0                0
    Unknown packets                   0                0
```

show pppoe underlying-interfaces extensive (PPPoE client in lockout condition)

```
user@host> show pppoe underlying-interfaces ge-1/0/0.0 extensive
ge-1/0/0.0 Index 71
  State: Static, Dynamic Profile: None,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Active Sessions: 0,
```



```
  Service Name Table: None,
  Direct Connect: Off,
  AC name: winona, Duplicate Protection: On,
  Short Cycle Protection: Off

  PacketType                        Sent         Received
    PADI                               0                7
    PADO                               3                0
    PADR                               0                3
    PADS                               3                0
    PADT                               2                1
    Service name error                 0                0
    AC system error                    0                0
    Generic error                      0                0
    Malformed packets                  0                0
    Unknown packets                    0                0

  Lockout Time (sec):  Min: 1, Max: 30
    Total clients in lockout: 1
    Total clients in lockout grace period: 0

  Client Address              Current     Elapsed    Next
    00:00:5e:00:53:11               4           3       8
```

show pppoe underlying-interfaces lockout

```
user@host> show pppoe underlying-interfaces ge-1/0/0.0 lockout
ge-1/0/0.0 Index 71
  Short Cycle Protection: Off,
  Lockout Time (sec):  Min: 10, Max: 60
    Total clients in lockout: 0
    Total clients in lockout grace period: 0
```

show pppoe underlying-interfaces detail (Autosensing Configured for ACI-based Dynamic VLANs)

```
user@host> show pppoe underlying-interfaces demux0.1073741826 detail
demux0.1073741826 Index 345
  State: Dynamic, Dynamic Profile: aci-vlan-pppoe-profile,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Active Sessions: 1,
  Agent Circuit Identifier: Autosensing,
  Service Name Table: None,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc,
  Short Cycle Protection: circuit-id,
```

show pppoe underlying-interfaces detail (Autosensing Configured for ALI-based Dynamic VLANs)

```
user@host> show pppoe underlying-interfaces demux0.1073741826 detail
demux0.1073741826 Index 345
  State: Dynamic, Dynamic Profile: aci-vlan-pppoe-profile,
  Max Sessions: 32000, Max Sessions VSA Ignore: Off,
  Active Sessions: 1,
  Line Identity: Autosensing,
  Service Name Table: None,
  Duplicate Protection: On, Short Cycle Protection: Off,
  Direct Connect: Off,
  AC Name: nbc,
  Short Cycle Protection: circuit-id,
```


show pppoe version

Syntax show pppoe version

Release Information Command introduced before Junos OS Release 7.4.

Description (M120 routers and M320 routers only) Display version information about PPPoE.

Options This command has no options.

Required Privilege Level view

List of Sample Output show pppoe version on page 2128

Output Fields Table 169 on page 2128 lists the output fields for the **show pppoe version** command. Output fields are listed in the approximate order in which they appear.

Table 169: show pppoe version Output Fields

Field Name	Field Description
version *n*	PPPoE version number and RFC. For example, **version 1, rfc 2516**.
PPPoE protocol	State of the PPPoE protocol: **enabled** or **disabled**.
Maximum Sessions	Maximum active sessions supported per router. The default is 256 sessions.
PADI resend timeout	Initial time, in seconds, that the router waits to receive a PPoE Active Discovery Offer (PADO) packet for the PPoE Active Discovery Initiation (PADI) packet sent. This timeout doubles for each successive PADI packet sent. Not displayed for M120 and M320 routers.
PADR resend timeout	Initial time, in seconds, that the router waits to receive a PPoE Active Discovery Session Confirmation (PADS) packet for the PPoE Active Discovery Request (PADR) packet sent. This timeout doubles for each successive PADR packet sent. Not displayed for M120 and M320 routers.
Max resend timeout	Maximum value, in seconds, that the PADI or PADR resend timer can accept. The maximum value is 64. Not displayed for M120 and M320 routers.
Max Configured AC timeout	Time, in seconds, during which the configured access concentrator must respond. Not displayed for M120 and M320 routers.

Sample Output

show pppoe version

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
   PPPoE protocol               = Enabled
   Maximum Sessions             = 256
```



```
PADI resend timeout          = 2 seconds
PADR resend timeout          = 16 seconds
Max resend timeout           = 64 seconds
Max Configured AC timeout    = 4 seconds
```

show protection-group ethernet-ring aps

Syntax **show protection-group ethernet-ring aps**

Release Information Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

Description Display the status of the Automatic Protection Switching (APS) and Ring APS (RAPS) messages on an Ethernet ring.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring interface on page 2145
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring statistics on page 2153
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output

show protection-group ethernet-ring aps (EX Switches) on page 2131
show protection-group ethernet-ring aps (Owner Node, Normal Operation on ACX and MX Routers) on page 2131
show protection-group ethernet-ring aps detail (Owner Node, Normal Operation on ACX and MX Routers) on page 2132
show protection-group ethernet-ring aps (MX RPL Owner Ring Node, Failure condition on non-RPL link of the ring) on page 2132
show protection-group ethernet-ring aps (MX Interconnection Ring Node, Failure condition in major ring on non-RPL link of the ring) on page 2132
show protection-group ethernet-ring aps (MX Series router) on page 2132
show protection-group ethernet-ring aps detail (MX Series router) on page 2132
show protection-group ethernet-ring aps (MX Interconnection Ring Node as RPL owner of major ring, rings in IDLE state) on page 2133
show protection-group ethernet-ring aps detail (EX2300 and EX3400 Switches) on page 2133

Output Fields Table 170 on page 2131 lists the output fields for the **show protection-group ethernet-ring aps** command. Output fields are listed in the approximate order in which they appear.

Table 170: show protection-group ethernet-ring aps Output Fields

Field Name	Field Description
Ethernet Ring	Name configured for the Ethernet ring.
Request/State	Status of the Ethernet ring RAPS messages. • **NR**—Indicates that there is no request for APS on the ring. • **SF**—Indicates that there is a signal failure on the ring. • **FS**—Indicates that there are active forced-switch requests in the ring. • **MS**—Indicates that there are active manual-switch requests in the ring. NOTE: Both **FS** and **MS** values are valid only when G.8032v2 is supported.
Ring Protection Link Blocked	Blocking on the ring protection link: **Yes** or **No**.
No Flush	Indicates the value of the Do Not Flush (DNF) flag in the received RAPS PDU. If the value is Yes, then FDB flush is not triggered as part of processing of the received RAPS PDU.
Blocked Port Reference	This parameter is the reference to the blocked ring port. If the east ring port is blocked, the Blocked Port Reference (BPR) value is 0. If the west ring port is blocked, the BPR value is 1. If both ring ports are blocked, this parameter can take any value. If both east and west ports are blocked or not blocked, the value would be 0. This field is valid only when G.8032v2 is supported.
Blocked Port Reference	Reference of the ring port on which traffic is blocked.
Originator	Indicates whether the node is the originator of the RAPS messages.
Remote Node ID	Identifier (in MAC address format) of the remote node.

Sample Output

show protection-group ethernet-ring aps (EX Switches)

```
user@switch>show protection-group ethernet-ring aps
Ring Name     Request/state  No Flush  RPL Blocked  Originator  Remote Node ID
erp1          NR             No        Yes          No          00:1F:12:30:B8:81
```

Sample Output

show protection-group ethernet-ring aps (Owner Node, Normal Operation on ACX and MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring   Request/state  RPL Blocked  No Flush  BPR  Originator  Remote
Node ID
Erp_1           NR             Yes          No        1    No
00:00:00:02:00:01
```


Sample Output

show protection-group ethernet-ring aps detail (Owner Node, Normal Operation on ACX and MX Routers)

```
user@host> show protection-group ethernet-ring aps detail
Ethernet-Ring name              : Erp_1
Request/State                   : NR
Ring Protection Link blocked    : Yes
No Flush Flag                   : No
Blocked Port Reference          : 1
Originator                      : No
Remote Node ID                  : 00:00:00:02:00:01
```

show protection-group ethernet-ring aps (MX RPL Owner Ring Node, Failure condition on non-RPL link of the ring)

```
user@host>show protection-group ethernet-ring aps
Ethernet Ring    Request/state    RPL Blocked  No Flush
pg101            SF               No           No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

show protection-group ethernet-ring aps (MX Interconnection Ring Node, Failure condition in major ring on non-RPL link of the ring)

```
user@host>show protection-group ethernet-ring aps
Ethernet Ring    Request/state  RPL Blocked  No Flush   BPR
pg_major         SF             No           No         0
pg_subring       NR             Yes          Yes        0

Originator  Remote Node ID
No          00:01:00:00:00:01
No          00:02:00:00:00:02
```

show protection-group ethernet-ring aps (MX Series router)

```
user@host>show protection-group ethernet-ring aps
Ethernet Ring    Request/state  RPL Blocked  No Flush   BPR  Originator  Remote
Node ID
Inst_Vlans_1-15 NR              Yes          Yes        1    Yes         NA

Inst_Vlans_16-30 NR             Yes          Yes        0    No
00:00:00:03:00:02
```

show protection-group ethernet-ring aps detail (MX Series router)

```
user@host>show protection-group ethernet-ring aps
Ethernet-Ring name              : Inst_Vlans_1-15
Request/State                   : NR
Ring Protection Link blocked    : Yes
No Flush Flag                   : Yes
Blocked Port Reference          : 1
Originator                      : Yes
Remote Node ID                  : NA

Ethernet-Ring name              : Inst_Vlans_16-30
Request/State                   : NR
Ring Protection Link blocked    : Yes
```



```
No Flush Flag                 : Yes
Blocked Port Reference        : 0
Originator                    : No
Remote Node ID                : 00:00:00:03:00:02
```

show protection-group ethernet-ring aps (MX Interconnection Ring Node as RPL owner of major ring, rings in IDLE state)

```
user@host>show protection-group ethernet-ring aps detail

Ethernet-Ring name            : pg_major
Request/State                 : NR
Ring Protection Link blocked  : Yes
No Flush Flag                 : Yes
Blocked Port Reference        : 0
Originator                    : Yes
Remote Node ID                : NA

Ethernet-Ring name            : pg_subring
Request/State                 : NR
Ring Protection Link blocked  : Yes
No Flush Flag                 : Yes
Blocked Port Reference        : 0
Originator                    : No
Remote Node ID                : 00:00:03:00:00:03
```

show protection-group ethernet-ring aps detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring aps detail
Ethernet-Ring name            : pg1001
Request/State                 : NR
Ring Protection Link blocked  : Yes
No Flush Flag                 : Yes
Blocked Port Reference        : 0
Originator                    : Yes
Remote Node ID                : NA
```

show protection-group ethernet-ring configuration

Syntax

show protection-group ethernet-ring configuration

Release Information

Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 14.1 for MX Series routers.
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

Description

Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring interface on page 2145
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring statistics on page 2153
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output

show protection-group ethernet-ring configuration (EX Switch) on page 2136
show protection-group ethernet-ring configuration detail (MX Series Router) on page 2137
show protection-group ethernet-ring configuration (MX Series Router) on page 2137
show protection-group ethernet-ring configuration detail (MX Series Router) on page 2137
show protection-group ethernet-ring configuration detail (MX Series Router) on page 2138
show protection-group ethernet-ring configuration (MX Series Router) on page 2138
show protection-group ethernet-ring configuration detail (MX Series Router) on page 2139

Output Fields

Table 171 on page 2134 lists the output fields for the **show protection-group ethernet-ring configuration** command. Output fields are listed in the approximate order in which they appear.

Table 171: show protection-group ethernet-ring configuration Output Fields

Output Fields	Field Description
G8032 Compatability Version	This is the compatibility version mode of ERP. This parameter always takes the value 1 in the case of G8032v1. This parameter is valid only for MX Series routers.
East Interface	One of the two switch interfaces that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 0.
West Interface	One of the two interfaces in a switch that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 1.

Table 171: show protection-group ethernet-ring configuration Output Fields (continued)

Output Fields	Field Description
Restore Interval	Configured interval of wait time after a link is restored. When a link goes down, the RPL link is activated. When the down link becomes active again, the RPL owner receives a notification. The RPL owner waits for the restore interval before issuing a block on the RPL link. The configured restore interval can be 5 through 12 minutes for ERPv1 and 1 through 12 minutes for ERPv2. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value. NOTE: Wait to Restore (WTR) configuration values on EX2300 and EX3400 switches must be 5-12 minutes.
Wait to Block Interval	Configured interval of wait time for link restoration when a manual command (manual switch or force switch) is cleared. On clearing the manual command, the RPL owner receives NR messages, which starts a timer with interval 'Wait to Block' to restore the RPL link after its expiration. This delay timer is set to be 5 seconds longer than the guard timer. The configured number can be from 5 seconds through 10 seconds. The parameter is valid only for G.8032v2. NOTE: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.
Guard Interval	Configured number of milliseconds (in 10 millisecond intervals, 10 milliseconds through 2000 milliseconds) that the node does not process any Ethernet ring protection protocol data units (PDUs). This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
Hold off interval	This is the interval at which the link is held down even before declaring that the link is down. Because the parameter is not supported at present, its value is always considered 0. This parameter is valid only for MX Series routers.
Node ID	Node ID for the switch or router. If the node ID is not configured, it is assigned by default. For EX Series switches, the Node ID value cannot be configured, whereas for MX Series routers, it can be configured.
Ring ID	In G8032v2, the ring ID can be within the range 1–239. All the nodes in a ring should have the same ring ID. In the case of G8032v1, the value of the ring ID is always 1. This parameter is valid only for MX Series routers.
Node Role	Indicates whether the ring node is operating as a normal ring-node or RPL-owner or RPL-neighbor. For G8032v1 RPL-neighbor role is not supported. This parameter is valid only for MX Series routers.

Table 171: show protection-group ethernet-ring configuration Output Fields (continued)

Output Fields	Field Description
Revertive Mode of Operation	This parameter indicates whether the ring is operating in revertive mode or nonrevertive mode. In nonrevertive mode of operation, when all links in the ring and Ethernet Ring Nodes have recovered and no external requests are active, the Ethernet Ring does not automatically revert. G8032v1 supports only revertive mode of operation. This parameter is valid only for MX Series routers.
RAPS Tx Dot1p priority	The RAPS Tx Dot1p priority is a parameter with which the RAPS is transmitted from the ring node. For G8032v1, the value of this parameter is always 0. For G8032v2, the value of this parameter can be within the range 0–7. This parameter is valid only for MX Series routers.
Node type	Indicates whether ring node is a normal ring node having two ring-links or a open ring-node having only a single ring-link or a interconnection ring-node. An interconnection ring node can be connected to major ring in non virtual-channel mode or in virtual channel mode. Ring interconnection is not supported for G8032v1. This parameter is valid only for MX Series routers.
Major ring name	If the node type is interconnection in the ring, this parameter takes the name of the major ring to which the sub-ring node is connected. This parameter is valid only for MX Series routers.
Interconnection mode	Indicates the interconnection mode if the type of the node is interconnection. An interconnection ring node can be connected to major ring in non-virtual channel mode or in virtual channel mode. This parameter is valid only for MX Series routers.
Propagate Topology Change event	When Propagate Topology Change event is set to 1, the change in the topology of sub-ring is propagated to the major ring, enabling the transmission of EVENT FLUSH RAPS PDU in the major ring. When the parameter is set to 0, the topology change in the sub-ring is not propagated to the major ring blocking EVENT FLUSH RAPS PDU transmission in the major ring. This parameter is valid only for MX Series routers.
Control Vlan	The VLAN that transfers ERP PDUs from one node to another.
Physical Ring	Physical ring if the east and west interfaces are nontrunk ports. For MX Series routers, the ring is termed a physical ring if no data channels are defined for the ring and the entire physical port forwarding is controlled by ERP.
Data Channel VLAN(s)	Data VLANs for which forwarding behavior is controlled by the ring instance.

Sample Output

show protection-group ethernet-ring configuration (EX Switch)

```
user@switch>show protection-group ethernet-ring configuration
Ethernet ring configuration parameters for protection group erp1
East Interface   : ge-0/0/3.0
West Interface   : ge-0/0/9.0
Restore Interval : 5 minutes
Guard Interval   : 500 ms
Node Id          : 00:1F:12:30:B8:81
```

```
Control Vlan     : 101
Physical Ring    : yes
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                    : 2
East interface (interface 0)                   : xe-2/3/0.1
West interface (interface 1)                   : xe-2/2/1.1
Restore interval                               : 5 minutes
Wait to Block interval                         : 5 seconds
Guard interval                                 : 500 ms
Hold off interval                              : 0 ms
Node ID                                        : 64:87:88:65:37:D0
Ring ID (1 ... 239)                            : 1
Node role (normal/rpl-owner/rpl-neighbour)     : normal
Revertive mode of operation                    : 1
RAPS Tx Dot1p priority (0 .. 7)                : 0
Node type (normal/open/interconnection)        : Normal
Control Vlan                                   : 100
Physical Ring                                  : No
Data Channel Vlan(s)                           : 200,300
```

show protection-group ethernet-ring configuration (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                    : 2
East interface (interface 0)                   : xe-2/3/0.1
West interface (interface 1)                   : xe-2/2/1.1
Restore interval                               : 5 minutes
Wait to Block interval                         : 5 seconds
Guard interval                                 : 500 ms
Hold off interval                              : 0 ms
Node ID                                        : 64:87:88:65:37:D0
Ring ID (1 ... 239)                            : 1
Node role (normal/rpl-owner/rpl-neighbour)     : rpl-neighbour
Node RPL end                                   : east-port
Revertive mode of operation                    : 1
RAPS Tx Dot1p priority (0 .. 7)                : 0
Node type (normal/open/interconnection)        : Normal
Control Vlan                                   : 100
Physical Ring                                  : No
Data Channel Vlan(s)                           : 200,300
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                    : 2
East interface (interface 0)                   : xe-2/3/0.1
West interface (interface 1)                   : xe-2/2/1.1
Restore interval                               : 5 minutes
Wait to Block interval                         : 5 seconds
Guard interval                                 : 500 ms
Hold off interval                              : 0 ms
Node ID                                        : 64:87:88:65:37:D0
Ring ID (1 ... 239)                            : 1
Node role (normal/rpl-owner/rpl-neighbour)     : rpl-owner
```



```
Node RPL end                                  : east-port
Revertive mode of operation                   : 1
RAPS Tx Dot1p priority (0 .. 7)               : 0
Node type (normal/open/interconnection)       : Normal
Control Vlan                                  : 100
Physical Ring                                 : No
Data Channel Vlan(s)                          : 200,300
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version                   : 2
East interface (interface 0)                  : xe-2/3/0.1
West interface (interface 1)                  : (no erp)
Restore interval                              : 5 minutes
Wait to Block interval                        : 5 seconds
Guard interval                                : 500 ms
Hold off interval                             : 0 ms
Node ID                                       : 64:87:88:65:37:D0
Ring ID (1 ... 239)                           : 1
Node role (normal/rpl-owner/rpl-neighbour)    : rpl-owner
Node RPL end                                  : east-port
Revertive mode of operation                   : 1
RAPS Tx Dot1p priority (0 .. 7)               : 0
Node type (normal/open/interconnection)       : Open
Control Vlan                                  : 100
Physical Ring                                 : No
Data Channel Vlan(s)                          : 200,300
```

show protection-group ethernet-ring configuration (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration
Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version                   : 2
East interface (interface 0)                  : xe-2/3/0.1
West interface (interface 1)                  : xe-2/2/1.1
Restore interval                              : 5 minutes
Wait to Block interval                        : 5 seconds
Guard interval                                : 500 ms
Hold off interval                             : 0 ms
Node ID                                       : 64:87:88:65:37:D0
Ring ID (1 ... 239)                           : 1
Node role (normal/rpl-owner/rpl-neighbour)    : rpl-owner
Node RPL end                                  : east-port
Revertive mode of operation                   : 1
RAPS Tx Dot1p priority (0 .. 7)               : 0
Node type (normal/open/interconnection)       : Normal
Control Vlan                                  : 100
Physical Ring                                 : No
Data Channel Vlan(s)                          : 200,300

Ethernet Ring configuration information for protection group pg_subring
G8032 Compatibility Version                   : 2
East interface (interface 0)                  : ge-2/0/0.1
West interface (interface 1)                  : (no erp)
Restore interval                              : 5 minutes
Wait to Block interval                        : 5 seconds
Guard interval                                : 500 ms
Hold off interval                             : 0 ms
Node ID                                       : 64:87:88:65:37:D0
```

```
Ring ID (1 ... 239)                       : 2
Node role (normal/rpl-owner/rpl-neighbour)  : normal
Revertive mode of operation               : 1
RAPS Tx Dot1p priority (0 .. 7)           : 0
Node type (normal/open/interconnection)   : Non-VC-Interconnection
Major ring name                           : pg_major
Interconnection mode (VC/Non-VC)          : Non-VC mode
Propagate Topology Change event           : 0
Control Vlan                              : 101
Physical Ring                             : No
Data Channel Vlan(s)                      : 200,300
```

show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version               : 2
East interface (interface 0)              : xe-2/3/0.1
West interface (interface 1)              : xe-2/2/1.1
Restore interval                          : 5 minutes
Wait to Block interval                    : 5 seconds
Guard interval                            : 500 ms
Hold off interval                         : 0 ms
Node ID                                   : 64:87:88:65:37:D0
Ring ID (1 ... 239)                       : 1
Node role (normal/rpl-owner/rpl-neighbour)  : rpl-owner
Node RPL end                              : east-port
Revertive mode of operation               : 1
RAPS Tx Dot1p priority (0 .. 7)           : 0
Node type (normal/open/interconnection)   : Normal
Control Vlan                              : 100
Physical Ring                             : No
Data Channel Vlan(s)                      : 200,300

Ethernet Ring configuration information for protection group pg_subring
G8032 Compatibility Version               : 2
East interface (interface 0)              : ge-2/0/0.1
West interface (interface 1)              : (no erp)
Restore interval                          : 5 minutes
Wait to Block interval                    : 5 seconds
Guard interval                            : 500 ms
Hold off interval                         : 0 ms
Node ID                                   : 64:87:88:65:37:D0
Ring ID (1 ... 239)                       : 2
Node role (normal/rpl-owner/rpl-neighbour)  : normal
Revertive mode of operation               : 1
RAPS Tx Dot1p priority (0 .. 7)           : 0
Node type (normal/open/interconnection)   : Non-VC-Interconnection
Major ring name                           : pg_major
Interconnection mode (VC/Non-VC)          : Non-VC mode
Propagate Topology Change event           : 0
Control Vlan                              : 101
Physical Ring                             : No
Data Channel Vlan(s)                      : 200,300
```

show protection-group ethernet-ring data-channel

Syntax

show protection-group ethernet-ring data-channel
<brief | detail>
<group-name *group-name*>

Release Information

Command introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

Description

Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.

Options

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

group-name—(Optional) Protection group for which to display statistics. If you omit this optional field, all protection group statistics for configured groups will be displayed.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring interface on page 2145
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring statistics on page 2153
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output

show protection-group ethernet-ring data-channel on page 2141
show protection-group ethernet-ring data-channel detail on page 2141
show protection-group ethernet-ring data-channel detail (EX2300 and EX3400 Switches) on page 2142

Output Fields

Table 172 on page 2140 lists the output fields for the **show protection-group ethernet-ring data-channel** command. Output fields are listed in the approximate order in which they appear.

Table 172: show protection-group ethernet-ring data-channel Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet ring.

Table 172: show protection-group ethernet-ring data-channel Output Fields (continued)

Field Name	Field Description
STP index	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.
Forward State	Forwarding state on the Ethernet ring. • **fowarding**—Indicates packets are being forwarded. • **discarding**—Indicates packets are being discarded.

Sample Output

show protection-group ethernet-ring data-channel

```
user@host> show protection-group ethernet-ring data-channel
Ethernet ring data channel information for protection group pg301

Interface     STP index  Forward State
xe-5/0/2      78         forwarding
xe-2/2/0      79         discarding

Ethernet ring data channel parameters for protection group pg302

Interface     STP index  Forward State
xe-5/0/2      80         forwarding
xe-2/2/0      81         forwarding
```

show protection-group ethernet-ring data-channel detail

```
user@host> show protection-group ethernet-ring data-channel detail
Ethernet ring data channel parameters for protection group pg301

Interface name                : xe-5/0/2
STP index                     : 78
Forward State                 : forwarding

Interface name                : xe-2/2/0
STP index                     : 79
Forward State                 : discarding


Ethernet ring data channel parameters for protection group pg302

Interface name                : xe-5/0/2
STP index                     : 80
Forward State                 : forwarding

Interface name                : xe-2/2/0
STP index                     : 81
Forward State                 : forwarding
```

show protection-group ethernet-ring data-channel detail (EX2300 and EX3400 Switches)

```
user@switch> show protection-group ethernet-ring data-channel detail
Ethernet ring data channel parameters for protection group pg1001

Interface name                : ge-0/0/42
STP index                     : 52
Forward State                 : discarding

Interface name                : ge-0/0/38
STP index                     : 53
Forward State                 : forwarding
```

show protection-group ethernet-ring flush-info

Syntax **show protection-group ethernet-ring flush-info**

Release Information Command introduced in Junos OS Release 14.2.

Description Display information about flush ports in an Ethernet ring.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring statistics on page 2153
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output show protection-group ethernet-ring flush-info (ACX and MX Series Routers) on page 2143
show protection-group ethernet-ring flush-info detail (ACX and MX Series Routers) on page 2144

Output Fields Table 173 on page 2143 lists the output fields for the **show protection-group ethernet-ring flush-info** command. Output fields are listed in the approximate order in which they appear.

Table 173: show protection-group ethernet-ring flush-info Output Fields

Field Name	Field Description
Interface	Physical interface configured for the Ethernet ring. This can be an aggregated Ethernet link also.
Originating Node	Node from which RAPS protocol data units originates on the Ethernet Ring.
Blocked Port Reference	Reference of the ring port on which traffic is blocked.

Sample Output

show protection-group ethernet-ring flush-info (ACX and MX Series Routers)

```
user@host> show protection-group ethernet-ring flush-info
```



```
Ethernet ring flush port information for protection group pg100

Interface     Originating Node   Blocked Port Reference
xe-5/0/2.4001 00:00:00:00:00:00   0
xe-2/2/0.4001 00:00:00:00:00:00   0
```

show protection-group ethernet-ring flush-info detail (ACX and MX Series Routers)

```
user@host> show protection-group ethernet-ring flush-info detail
Ethernet ring flush port information for protection group pg100

Interface name                : xe-5/0/2.4001
Originating Node              : 00:00:00:00:00:00
Blocked Port Reference        : 0

Interface name                : xe-2/2/0.4001
Originating Node              : 00:00:00:00:00:00
Blocked Port Reference        : 0
```

show protection-group ethernet-ring interface

Syntax
show protection-group ethernet-ring interface

Release Information
Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.3X54 for ACX Series routers.
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

Description
Displays the status of the Automatic Protection Switching (APS) interfaces on an Ethernet ring.

Options
This command has no options.

Required Privilege Level
view

Related Documentation

- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring statistics on page 2153
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output

Output Fields
Table 174 on page 2146 lists the output fields for both the EX Series switch, and the ACX Series and MX Series router **show protection-group ethernet-ring interface** commands. Output fields are listed in the approximate order in which they appear.

Table 174: MX Series Routers show protection-group ethernet-ring interface Output Fields

Field Name	Field Description
Ethernet ring port parameters for protection group *group-name*	Output is organized by configured protection group.
Interface	Physical interfaces configured for the Ethernet ring. This can be an aggregated Ethernet link also.
Control Channel	(MX Series router only) Logical unit configured on the physical interface.
Direction	Direction of the traffic.
Forward State	State of the ring forwarding on the interface: **discarding** or **forwarding**.
Ring Protection Link End	Whether this interface is the end of the ring: **Yes** or **No**.
Signal Failure	Whether there a signal failure exists on the link: **Clear** or **Set**.
Admin State	State of the interface: For EX switches, **ready**, **ifl ready**, or **waiting**. For MX routers, **IFF ready** or **IFF disabled**.

Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface      Forward State  RPL End  Signal Failure  Admin State

ge-0/0/3.0     discarding     Yes        Clear           ready
ge-0/0/9.0     forwarding     No         Clear           ready
```

show protection-group ethernet-ring interface (Owner Node MX Series Router)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface  Control Channel  Direction  Forward State  RPL End  SF     Admin State
ge-1/2/0   ge-1/2/0.100     east       forwarding     No        Clear  IFF ready
ge-1/2/2   ge-1/2/2.100     west       forwarding     No        Clear  IFF ready
```

show protection-group ethernet-ring interface detail (Owner Node MX Series Router)

```
user@host> show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group pg101

Interface name                : ge-1/2/0
Control channel name          : ge-1/2/0.100
```



```
Interface direction            : east
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready

Interface name                 : ge-1/2/2
Control channel name           : ge-1/2/2.100
Interface direction            : west
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready
```

show protection-group ethernet-ring interface (EX Series Switch Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Ethernet ring port parameters for protection group pg101

Interface      Forward State  RPL End  Signal Failure  Admin State

ge-0/0/3.0     discarding     Yes        Clear          ready
ge-0/0/9.0     forwarding     No         Clear          ready
```

show protection-group ethernet-ring interface detail (ACX Series and MX Series)

```
user@host> show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group Erp_1

Interface name                 : xe-0/0/0
Control channel name           : xe-0/0/0.1
Interface direction            : east
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready

Interface name                 : et-0/0/48
Control channel name           : et-0/0/48.1
Interface direction            : west
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready
```

show protection-group ethernet-ring interface detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group pg1001

Interface name                 : ge-0/0/14
Control channel name           : ge-0/0/14.0
Interface direction            : east
Ring Protection Link End       : No
Signal Failure                 : Clear
Forward State                  : forwarding
Interface Admin State          : IFF ready
```



```
Interface name                : ge-0/0/18
Control channel name          : ge-0/0/18.0
Interface direction           : west
Ring Protection Link End      : No
Signal Failure                : Clear
Forward State                 : forwarding
Interface Admin State         : IFF ready
```

show protection-group ethernet-ring interface detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring interface detail
Ethernet ring port parameters for protection group pg1001

Interface name                : ge-0/0/42
Control channel name          : ge-0/0/42.0
Interface direction           : east
Ring Protection Link End      : Yes
Signal Failure                : Clear
Forward State                 : discarding
Interface Admin State         : IFF ready

Interface name                : ge-0/0/38
Control channel name          : ge-0/0/38.0
Interface direction           : west
Ring Protection Link End      : No
Signal Failure                : Clear
Forward State                 : forwarding
Interface Admin State         : IFF ready
```


show protection-group ethernet-ring node-state

Syntax **show protection-group ethernet-ring node-state**

Release Information Command introduced in Junos OS Release 9.4 for MX Series routers.
Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 12.3X54 for ACX Series routers.
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

Description Display the status of the Automatic Protection Switching (APS) nodes on an Ethernet ring.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring interface on page 2145
- show protection-group ethernet-ring statistics on page 2153
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output

show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Normal Operation) on page 2151
show protection-group ethernet-ring node-state (MX Series Router - Normal Ring Node, Normal Operation) on page 2151
show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Remote Failure Condition) on page 2151
show protection-group ethernet-ring node-state detail (ACX Series and MX Series Router) on page 2151
show protection-group ethernet-ring node-state detail (MX Series Router - RPL Owner Node, Normal Operation) on page 2152
show protection-group ethernet-ring node-state detail (EX2300 and EX3400 Switches) on page 2152

Output Fields Table 175 on page 2149 lists the output fields for the **show protection-group ethernet-ring node-state** command. Output fields are listed in the approximate order in which they appear.

Table 175: show protection-group ethernet-ring node-state Output Fields

Field Name	Field Description
Ring Name/Ethernet Ring	Name configured for the Ethernet ring.

Table 175: show protection-group ethernet-ring node-state Output Fields (continued)

Field Name	Field Description
APS State	State of the Ethernet ring APS. • **idle**—Indicates that the ring is working in normal condition and there is no active or pending protection-switching request in the ring. When the ring is in idle state, it is blocked at the RPL link. • **protected**—Indicates that there is a protection switch on the ring because of a signal failure condition on the ring link. • **MS**—Indicates that the manual switch command is active in the ring. • **FS**—Indicates that the forced switch command is active in the ring. • **pending**—Indicates that the ring is in pending state.
Event	Events on the ring. • **NR-RB**—Indicates that there is no APS request and the ring link is blocked on the ring owner node. • **NR**—Indicates that there is no APS request pending in the ring. • **local SF**—Indicates that there is signal failure on one or both of the ring links of the node. • **remote SF**—Indicates that there is signal failure on one or more ring links of any other node of the ring. • **local FS**—Indicates that there is a forced switched command active on one or both of the ring links of the node. • **remote FS**—Indicates that there is a forced switch command active on one or more ring links of any other node of the ring. • **local MS**—Indicates that there is a manual switch command active on one of the ring links of the node. • **remote MS**—Indicates that there is a manual switch command active on one or more ring links of any other node of the ring. • **WTR running**—Indicates that the wait to restore timer is running on the RPL owner. • **WTB running**—Indicates that the wait to block timer is running on the RPL owner.
RPL Owner / Ring Protection Link Owner	Whether this node is the ring owner: **Yes** or **No**.
WTR Timer / Restore Timer	Restoration timer: **running** or **disabled**.
WTB Timer / Wait to block timer	Wait to block timer: **running** or **disabled**. NOTE: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.

Table 175: show protection-group ethernet-ring node-state Output Fields (continued)

Field Name	Field Description
Wait to block timer (WTB Timer)	Wait to block interval. NOTE: The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.
Guard Timer	Guard timer: **running** or **disabled**.
Op State / Operational State	State of the node: **Operational** or **any internal wait state.**.

Sample Output

show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring   APS State   Event       RPL Owner  WTR Timer  WTB Timer  Guard
Timer  Operation state
pg101           idle        NR-RB       Yes        disabled   disabled   disabled
     operational
pg102            idle       NR-RB       No         disabled   disabled   disabled
     operational
```

show protection-group ethernet-ring node-state (MX Series Router - Normal Ring Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         RPL Owner
pg102            idle         NR-RB         No

WTR Timer  WTB Timer   Guard Timer  Operation state
disabled   disabled    disabled      operational
```

show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Remote Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event         RPL Owner
pg101            protected    remote SF         Yes

WTR Timer   WTB Timer Guard Timer  Operation state
disabled    disabled   disabled     operational
```

show protection-group ethernet-ring node-state detail (ACX Series and MX Series Router)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : Erp_1
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : No
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled
```



```
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring node-state detail (MX Series Router - RPL Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : pg101
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : Yes
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled
Guard Timer                     : disabled
Operation state                 : operational

Ethernet-Ring name              : pg102
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : No
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring node-state detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : pg1001
APS State                       : idle
Event                           : NR-RB
Ring Protection Link Owner      : Yes
Wait to Restore Timer           : disabled
Wait to Block Timer             : disabled    <-field not supported. Always
disabled.
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring statistics

Syntax

show protection-group ethernet-ring statistics group-name ***group-name***

<brief | detail>

Release Information

Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

Description

Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.

Options

group-name—Display statistics for the protection group. If you omit this option, protection group statistics for all configured groups are displayed.

brief—Display brief statistics for the protection group.

detail—Display detailed statistics for the protection group.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring interface on page 2145
- show protection-group ethernet-ring vlan on page 2159

List of Sample Output

show protection-group ethernet-ring statistics (EX Series Switch) on page 2155
show protection-group ethernet-ring statistics (MX Series Router) on page 2155
show protection-group ethernet-ring statistics detail (Specific Group)(MX Series Router) on page 2156
show protection-group ethernet-ring statistics (Owner Node, Failure Condition on ACX and MX Router) on page 2156
show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on ACX and MX Router) on page 2157
show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches) on page 2157
show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches) on page 2157

Output Fields

Table 176 on page 2154 lists the output fields for the **show protection-group ethernet-ring statistics** command.

Table 176: show protection-group ethernet-ring statistics Output Fields

Field Name	Field Description
Ethernet Ring Statistics for PG	Name of the protection group for which statistics are displayed.
RAPS event sent	Number of times Ring Automatic Protection Switching (RAPS) message transmission event occurred locally. This field is applicable only to MX Series routers.
RAPS event received	Number of RAPS messages received and processed by ERP state-machine and which resulted in state transition. This field is applicable only to MX Series routers.
Local SF	Number of times a signal failure has occurred locally.
Remote SF	Number of times a signal failure has occurred anywhere else on the ring.
NR event	Number of times a No Request event has occurred on the ring. This field is applicable only to EX Series switches.
NR event sent	Number of times a No Request event has occurred locally. This field is applicable only to MX Series routers.
NR event received	Number of times a No Request event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.
NR-RB event	Number of times a No Request, Ring Blocked event has occurred on the ring. This field is applicable only to EX Series switches.
NR-RB event sent	Number of times a No Request, Ring Blocked event has occurred locally. This field is applicable only to MX Series routers.
NR-RB event received	Number of times a No Request, Ring Blocked event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.
Flush event sent	Number of times flush-event RAPS message transmission event occurred locally. This field is applicable only to MX Series routers.
Flush event received	Number of flush-event RAPS messages received and processed by the ring instance control process. This field is applicable only to MX Series routers.
Local FS event sent	Number of times a forced switch event has occurred locally. This field is applicable only to MX Series routers.
Remote FS event received	Number of times a forced switch event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.
Local MS event sent	Number of times a manual switch event has occurred locally. This field is applicable only to MX Series routers.

Table 176: show protection-group ethernet-ring statistics Output Fields (continued)

Field Name	Field Description
Remote MS event received	Number of times a manual switch event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.

Table 177 on page 2155 lists the output fields for the **show protection-group ethernet-ring statistics** command when the **detail** option is used. These fields are valid only for MX Series routers.

Table 177: show protection-group ethernet-ring statistics detail Output Fields (for MX Series Routers)

Field Name	Field Description
Total number of FDB flush	Number of times forwarding database (FDB) flush has happened for the ring instance.
Flush-logic triggered flush	Number of times FDB flush has happened because of flush-logic based on node ID and Blocked Port Reference (BPR).
Remote RAPS PDU received	Number of valid RAPS PDU messages received. This counter counts only RAPS messages generated by other devices on the ring.
Remote RAPS dropped due to guard-timer	Number of RAPS messages dropped by the device because the guard timer is running.
Invalid remote RAPS PDU dropped	Number of RAPS messages dropped by the device because the messages are invalid.
RAPS dropped due to miscellaneous errors	Number of RAPS messages dropped because of any other reason. For example, messages dropped because of unsupported functionality.
Local received RAPS PDU dropped	Number of self-generated RAPS messages received and dropped.

Sample Output

show protection-group ethernet-ring statistics (EX Series Switch)

```
user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1       2       1       2        3
```

show protection-group ethernet-ring statistics (MX Series Router)

```
user@host> show protection-group ethernet-ring statistics
Ethernet Ring statistics for PG Pg-1
RAPS event sent                   : 1
RAPS event received               : 1152
Local SF happened:                : 0
Remote SF happened:               : 428
```



```
NR event sent:                  : 1
NR event received:              : 133
NR-RB event sent:               : 0
NR-RB event received:           : 591
Flush event sent                : 0
Flush event received:           : 0
Local FS event sent:            : 0
Remote FS event received:       : 0
Local MS event sent:            : 0
Remote MS event received:       : 0
```

show protection-group ethernet-ring statistics detail (Specific Group)(MX Series Router)

```
user@host> show protection-group ethernet-ring statistics detail
Ethernet Ring statistics for PG Pg-1
RAPS event sent                         : 1
RAPS event received                     : 0
Local SF happened                       : 0
Remote SF happened                      : 0
NR event sent                           : 1
NR event received                       : 0
NR-RB event sent                        : 0
NR-RB event received                    : 0
Flush event sent                        : 0
Flush event received                    : 0
Local FS event sent                     : 0
Remote FS event received                : 0
Local MS event sent                     : 0
Remote MS event received                : 0
Total number of FDB flush               : 0
Flush-logic triggered flush             : 0
Remote raps PDU received                : 0
Remote raps dropped due to guard-timer  : 0
Invalid remote raps PDU dropped         : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped         : 0
```

show protection-group ethernet-ring statistics (Owner Node, Failure Condition on ACX and MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                       : 1
RAPS received                   : 0
Local SF happened:              : 0
Remote SF happened:             : 0
NR event happened:              : 0
NR-RB event happened:           : 1
NR event sent:                  : 0
NR event received:              : 0
NR-RB event sent:               : 1
NR-RB event received:           : 0
Flush event sent                : 0
Flush event received:           : 0
Local FS event sent:            : 0
Remote FS event received:       : 0
Local MS event sent:            : 0
Remote MS event received:       : 0
```

show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on ACX and MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                        : 1
RAPS received                    : 0
Local SF happened:               : 0
Remote SF happened:              : 0
NR event happened:               : 0
NR-RB event happened:            : 1
NR event sent:                   : 0
NR event received:               : 0
NR-RB event sent:                : 1
NR-RB event received:            : 0
Flush event sent                 : 0
Flush event received:            : 0
Local FS event sent:             : 0
Remote FS event received:        : 0
Local MS event sent:             : 0
Remote MS event received:        : 0
```

show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring statistics detail
Ethernet Ring statistics for PG pg1001
RAPS event sent                          : 1
RAPS event received                      : 1
Local SF happened                        : 0
Remote SF happened                       : 0
NR event sent                            : 1
NR event received                        : 0
NR-RB event sent                         : 0
NR-RB event received                     : 1
Flush event sent                         : 0
Flush event received                     : 0
Local FS event sent                      : 0
Remote FS event received                 : 0
Local MS event sent                      : 0
Remote MS event received                 : 0
Total number of FDB flush                : 0
Flush-logic triggered flush              : 0
Remote raps PDU received                 : 145
Remote raps dropped due to guard-timer   : 0
Invalid remote raps PDU dropped          : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped          : 0
```

show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring statistics detail
Ethernet Ring statistics for PG pg1001
RAPS event sent                          : 2
RAPS event received                      : 0
Local SF happened                        : 0
Remote SF happened                       : 0
NR event sent                            : 1
NR event received                        : 0
NR-RB event sent                         : 1
NR-RB event received                     : 0
Flush event sent                         : 0
```

```
Flush event received                     : 0
Total number of FDB flush                : 0
Remote raps PDU received                 : 211
Remote raps dropped due to guard-timer   : 0
Invalid remote raps PDU dropped          : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped          : 91
```

show protection-group ethernet-ring vlan

Syntax

show protection-group ethernet-ring vlan
<brief | detail>
<group-name *group-name*>

Release Information

Command introduced in Junos OS Release 10.2.
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

Description

On MX Series routers, display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

Options

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

group-name—(Optional) Protection group for which to display details such as data channel interfaces, vlan, and bridge-domain. If you omit this optional field, details for all configured protection groups will be displayed.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring aps on page 2130
- show protection-group ethernet-ring data-channel on page 2140
- show protection-group ethernet-ring interface on page 2145
- show protection-group ethernet-ring node-state on page 2149
- show protection-group ethernet-ring statistics on page 2153

List of Sample Output

show protection-group ethernet-ring vlan on page 2160
show protection-group ethernet-ring vlan brief on page 2160
show protection-group ethernet-ring vlan detail on page 2161
show protection-group ethernet-ring vlan group-name vkm01 on page 2162
show protection-group ethernet-ring vlan detail (EX2300 and EX3400 Switches) on page 2162

Output Fields

Table 178 on page 2159 lists the output fields for the **show protection-group ethernet-ring vlan** command. Output fields are listed in the approximate order in which they appear.

Table 178: show protection-group ethernet-ring vlan Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet protection ring.
Vlan	Name of the VLAN associated with the interface configured for the Ethernet protection ring.

Table 178: show protection-group ethernet-ring vlan Output Fields (continued)

Field Name	Field Description
STP index	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.
Bridge Domain	Name of the bridge domain that is associated with the VLAN configured for the Ethernet protection ring.

Sample Output

show protection-group ethernet-ring vlan

```
user@host> show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group vkm01

Interface     Vlan    STP Index    Bridge Domain
xe-5/0/2      1       78           default-switch/bd1
xe-2/2/0      1       79           default-switch/bd1
xe-5/0/2      2       78           default-switch/bd2
xe-2/2/0      2       79           default-switch/bd2
xe-5/0/2      3       78           default-switch/bd3
xe-2/2/0      3       79           default-switch/bd3
xe-5/0/2      4       78           default-switch/bd4
xe-2/2/0      4       79           default-switch/bd4
xe-5/0/2      5       78           default-switch/bd5
xe-2/2/0      5       79           default-switch/bd5
xe-5/0/2      6       78           default-switch/bd6
xe-2/2/0      6       79           default-switch/bd6
xe-5/0/2      7       78           default-switch/bd7
xe-2/2/0      7       79           default-switch/bd7
xe-5/0/2      8       78           default-switch/bd8
xe-2/2/0      8       79           default-switch/bd8
xe-5/0/2      9       78           default-switch/bd9
xe-2/2/0      9       79           default-switch/bd9
xe-5/0/2      10      78           default-switch/bd10
xe-2/2/0      10      79           default-switch/bd10
xe-5/0/2      11      78           default-switch/bd11
xe-2/2/0      11      79           default-switch/bd11
xe-5/0/2      12      78           default-switch/bd12
xe-2/2/0      12      79           default-switch/bd12
xe-5/0/2      13      78           default-switch/bd13
xe-2/2/0      13      79           default-switch/bd13
xe-5/0/2      14      78           default-switch/bd14
xe-2/2/0      14      79           default-switch/bd14
xe-5/0/2      15      78           default-switch/bd15
xe-2/2/0      15      79           default-switch/bd15
```

show protection-group ethernet-ring vlan brief

```
user@host> show protection-group ethernet-ring vlan brief
```

```
Ethernet ring IFBD parameters for protection group vkm01

Interface     Vlan     STP Index    Bridge Domain
xe-5/0/2      1        78           default-switch/bd1
xe-2/2/0      1        79           default-switch/bd1
xe-5/0/2      2        78           default-switch/bd2
xe-2/2/0      2        79           default-switch/bd2
xe-5/0/2      3        78           default-switch/bd3
xe-2/2/0      3        79           default-switch/bd3
xe-5/0/2      4        78           default-switch/bd4
xe-2/2/0      4        79           default-switch/bd4
xe-5/0/2      5        78           default-switch/bd5
xe-2/2/0      5        79           default-switch/bd5
xe-5/0/2      6        78           default-switch/bd6
xe-2/2/0      6        79           default-switch/bd6
xe-5/0/2      7        78           default-switch/bd7
xe-2/2/0      7        79           default-switch/bd7
xe-5/0/2      8        78           default-switch/bd8
xe-2/2/0      8        79           default-switch/bd8
xe-5/0/2      9        78           default-switch/bd9
xe-2/2/0      9        79           default-switch/bd9
xe-5/0/2      10       78           default-switch/bd10
xe-2/2/0      10       79           default-switch/bd10
xe-5/0/2      11       78           default-switch/bd11
xe-2/2/0      11       79           default-switch/bd11
xe-5/0/2      12       78           default-switch/bd12
xe-2/2/0      12       79           default-switch/bd12
xe-5/0/2      13       78           default-switch/bd13
xe-2/2/0      13       79           default-switch/bd13
xe-5/0/2      14       78           default-switch/bd14
xe-2/2/0      14       79           default-switch/bd14
xe-5/0/2      15       78           default-switch/bd15
xe-2/2/0      15       79           default-switch/bd15
```

show protection-group ethernet-ring vlan detail

```
user@host> show protection-group ethernet-ring vlan detail
Ethernet ring IFBD parameters for protection group vkm01

Interface name                : xe-5/0/2
Vlan                          : 1
STP index                     : 78
Bridge Domain                 : default-switch/bd1

Interface name                : xe-2/2/0
Vlan                          : 1
STP index                     : 79
Bridge Domain                 : default-switch/bd1

Interface name                : xe-5/0/2
Vlan                          : 2
STP index                     : 78
Bridge Domain                 : default-switch/bd2

Interface name                : xe-2/2/0
Vlan                          : 2
STP index                     : 79
Bridge Domain                 : default-switch/bd2

Interface name                : xe-5/0/2
Vlan                          : 3
```

```
STP index                    : 78
Bridge Domain                : default-switch/bd3
```

show protection-group ethernet-ring vlan group-name vkm01

```
user@host> show protection-group ethernet-ring vlan vkm01

Ethernet ring IFBD parameters for protection group vkm01

Interface     Vlan    STP Index    Bridge Domain
xe-5/0/2      16      80           default-switch/bd16
xe-2/2/0      16      81           default-switch/bd16
xe-5/0/2      17      80           default-switch/bd17
xe-2/2/0      17      81           default-switch/bd17
xe-5/0/2      18      80           default-switch/bd18
xe-2/2/0      18      81           default-switch/bd18
xe-5/0/2      19      80           default-switch/bd19
xe-2/2/0      19      81           default-switch/bd19
xe-5/0/2      20      80           default-switch/bd20
xe-2/2/0      20      81           default-switch/bd20
xe-5/0/2      21      80           default-switch/bd21
xe-2/2/0      21      81           default-switch/bd21
xe-5/0/2      22      80           default-switch/bd22
xe-2/2/0      22      81           default-switch/bd22
xe-5/0/2      23      80           default-switch/bd23
xe-2/2/0      23      81           default-switch/bd23
xe-5/0/2      24      80           default-switch/bd24
xe-2/2/0      24      81           default-switch/bd24
xe-5/0/2      25      80           default-switch/bd25
xe-2/2/0      25      81           default-switch/bd25
xe-5/0/2      26      80           default-switch/bd26
xe-2/2/0      26      81           default-switch/bd26
xe-5/0/2      27      80           default-switch/bd27
xe-2/2/0      27      81           default-switch/bd27
xe-5/0/2      28      80           default-switch/bd28
xe-2/2/0      28      81           default-switch/bd28
xe-5/0/2      29      80           default-switch/bd29
xe-2/2/0      29      81           default-switch/bd29
xe-5/0/2      30      80           default-switch/bd30
xe-2/2/0      30      81           default-switch/bd30
```

show protection-group ethernet-ring vlan detail (EX2300 and EX3400 Switches)

```
user@switch>show protection-group ethernet-ring vlan detail
Ethernet ring IFBD parameters for protection group pg1001

Interface name               : ge-0/0/42
Vlan                         : 2001
STP index                    : 52
Bridge Domain                : default-switch/vlan2001

Interface name               : ge-0/0/38
Vlan                         : 2001
STP index                    : 53
Bridge Domain                : default-switch/vlan2001
```


show security macsec connections (MX Series)

Syntax

show security macsec connections
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers.
Support for MPC7E-10G introduced in Junos OS Release 16.1R1 for MX240, MX480, and MX960 routers.

Description

Display the status of the active MACsec connections on the router.

Options

none—Display MACsec connection information for all interfaces on the switch.

interface ***interface-name***—(Optional) Display MACsec connection information for the specified interface only.

Required Privilege Level

view

List of Sample Output

show security macsec connections on page 2164
show security macsec connections (MX480 routers with MPC7E-10G) on page 2164
show security macsec connections (MX480 routers with MPC7E-10G) on page 2165

Output Fields

Table 179 on page 2163 lists the output fields for the **show security macsec connections** command. Output fields are listed in the approximate order in which they appear.

Table 179: show security macsec connections Output Fields

Field Name	Field Description
Fields for Interface	
Interface name	Name of the interface.
CA name	Name of the connectivity association. A connectivity association is named using the **connectivity-association** statement when you are enabling MACsec.
Cipher suite	Name of the cipher suite used for encryption.
Encryption	Encryption setting. Encryption is enabled when this output is **on** and disabled when this output is **off**. The encryption setting is set using the **no-encryption** statement in the connectivity association when using static connectivity association key (CAK) security mode and is set using the **encryption** statement in the secure channel when using static secure association key (SAK) or dynamic security mode.

Table 179: show security macsec connections Output Fields (continued)

Field Name	Field Description
Key server offset	The offset value in a packet from which encryption can be performed. The offset is set using the **offset** statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode or the secure channel when using static secure association key (SAK) security mode.
Include SCI	SCI tagging. The SCI tag is included on packets in a secure channel when this output is **yes**, and not included on packets in a secure channel when this output is **no**. SCI tagging is automatically enabled on MX Series routers. By default, include SCI tag is disabled. You can enable SCI tagging using the **include-sci** statement in the connectivity association configuration.
Replay protect	By default, replay protection is disabled. Replay protection ensures that a snooped packet is not replayed or a packet number is reused. Replay protection is enabled when this output is **on** and disabled when this output is **off**. You can enable replay protection using the **replay-protect** statement in the connectivity association configuration.
Replay window	Number of packets that can be replayed. Must be configured with replay protection. This output is set to **0** when replay protection is disabled, and is the size of the replay window, in number of packets, when replay protection is enabled. The size of the replay window is configured using the **replay-window-size** statement in the connectivity association configuration.

Sample Output

show security macsec connections

```
user@host> show security macsec connections
Interface name: xe-0/1/0
        CA name: CA1
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
```

show security macsec connections (MX480 routers with MPC7E-10G)

```
user@host> show security macsec connections
 Interface name: xe-4/0/18
        CA name: ca1
        Cipher suite: GCM-AES-128   Encryption: on
        Key server offset: 30       Include SCI: no
        Replay protect: off         Replay window: 0
          Outbound secure channels
            SC Id: 54:1E:56:B4:0D:3A/1
            Outgoing packet number: 11
            Secure associations
            AN: 1 Status: inuse Create time: 1d 17:31:10
          Inbound secure channels
            SC Id: 54:1E:56:B3:CA:A7/1
```

```
      Secure associations
      AN: 1 Status: inuse Create time: 1d 17:31:10
```

show security macsec connections (MX480 routers with MPC7E-10G)

```
user@host> show security macsec connections interface xe-1/0/7
CA name: caae1
        Cipher suite: AES_GCM_128   Encryption: off
        Key server offset: 0        Include SCI: no
        Replay protect: off         Replay window: 0
          Outbound secure channels
            SC Id: 54:1E:56:B3:CA:9C/1
            Outgoing packet number: 1
            Secure associations
            AN: 0 Status: inuse Create time: 4d 05:56:06
          Inbound secure channels
            SC Id: 54:1E:56:B4:0D:2F/1
            Secure associations
            AN: 0 Status: inuse Create time: 4d 05:56:06
```

show security macsec statistics (MX Series)

Syntax

show security macsec statistics
<brief | detail>
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers. Support for MPC7E-10G introduced in Junos OS Release 16.1R1 for MX240, MX480, and MX960 routers.

Description

Display Media Access Control Security (MACsec) statistics.

Options

none—Display MACsec statistics in brief form for all interfaces on the router.

brief | detail—(Optional) Display the specified level of output. Using the **brief** option is equivalent to entering the command with no options (the default). The **detail** option displays additional fields that are not visible in the **brief** output.

NOTE: The field names that only appear in this command output when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel.

interface interface-name—(Optional) Display MACsec statistics for the specified interface only.

Required Privilege Level

view

List of Sample Output

show security macsec statistics interface detail on page 2168
show security macsec statistics (MX480 router with MPC7E-10G) on page 2169
show security macsec statistics (MX480 router with MPC7E-10G) on page 2169
show security macsec statistics detail (MX480 router with MPC7E-10G) on page 2170

Output Fields

Table 180 on page 2166 lists the output fields for the **show security macsec statistics** command. Output fields are listed in the approximate order in which they appear.

The field names that appear in this command's output only when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel. Those field names are, therefore, not included in this table.

Table 180: show security macsec statistics Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface.	All levels

Table 180: show security macsec statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Fields for Secure Channel transmitted		
Encrypted packets	Total number of packets transmitted out of the interface in the secure channel that were secured and encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
Encrypted bytes	Total number of bytes transmitted out of the interface in the secure channel that were secured and encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
Protected packets	Total number of packets transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
Protected bytes	Total number of bytes transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec. Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).	All levels
Fields for Secure Association transmitted		
Encrypted packets	Total number of packets transmitted out of the interface in the connectivity association that were secured and encrypted using MACsec. The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).	All levels
Protected packets	Total number of packets transmitted out of the interface in the connectivity association that were secured but not encrypted using MACsec. The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).	All levels
Fields for Secure Channel received		
Accepted packets	The number of received packets that have been accepted by the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check. This counter increments for traffic that is and is not encrypted using MACsec.	All levels

Table 180: show security macsec statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Validated bytes	The number of bytes that have been validated by the MACsec integrity check and received on the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link. This counter does not increment when MACsec encryption is disabled.	All levels
Decrypted bytes	The number of bytes received in the secure channel on the interface that have been decrypted. The secure channel is used to send all data plane traffic on a MACsec-enabled link. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.	All levels
Fields for Secure Association received		
Accepted packets	The number of received packets that have been accepted in the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.	All levels
Validated bytes	The number of bytes that have been validated by the MACsec integrity check and received on the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface. This counter does not increment when MACsec encryption is disabled.	All levels
Decrypted bytes	The number of bytes received in the connectivity association on the interface that have been decrypted. The counter includes all control and data plane traffic accepted on the interface. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.	All levels

Sample Output

show security macsec statistics interface detail

```
user@host> show security macsec statistics interface xe-0/1/0 detail

Interface name: xe-0/1/0
  Secure Channel transmitted
      Encrypted packets: 123858
      Encrypted bytes:   32190903
      Protected packets: 0
      Protected bytes:   0
```

```
    Secure Association transmitted
        Encrypted packets: 123858
        Protected packets: 0
    Secure Channel received
        Accepted packets:  123877
        Validated bytes:   0
        Decrypted bytes:   32196238
    Secure Association received
        Accepted packets:  123877
        Validated bytes:   0
        Decrypted bytes:   32196238
    Error and debug
    Secure Channel transmitted packets
        Untagged: 0, Too long: 0
    Secure Channel received packets
        Control: 0, Tagged miss: 3202804
        Untagged hit: 0, Untagged: 0
        No tag: 0, Bad tag: 0
        Unknown SCI: 0, No SCI: 0
        Control pass: 0, Control drop: 0
        Uncontrol pass: 123877, Uncontrol drop: 0
        Hit dropped: 0, Invalid accept: 0
        Late drop: 0, Delayed accept: 0
        Unchecked: 0, Not valid drop: 0
        Not using SA drop: 0, Unused SA accept: 0
```

show security macsec statistics (MX480 router with MPC7E-10G)

```
user@host> show security macsec statistics
Interface name: xe-4/0/18
    Secure Channel transmitted
        Encrypted packets: 10
        Encrypted bytes:   840
        Protected packets: 0
        Protected bytes:   0
    Secure Association transmitted
        Encrypted packets: 10
        Protected packets: 0
    Secure Channel received
        Accepted packets:  0
        Validated bytes:   0
        Decrypted bytes:   0
    Secure Association received
        Accepted packets:  0
        Validated bytes:   0
        Decrypted bytes:   0
```

show security macsec statistics (MX480 router with MPC7E-10G)

```
user@host> show security macsec statistics interface xe-1/0/7
Secure Channel transmitted
        Encrypted packets: 0
        Encrypted bytes:   0
        Protected packets: 0
        Protected bytes:   0
    Secure Association transmitted
        Encrypted packets: 0
        Protected packets: 0
    Secure Channel received
        Accepted packets:  0
```

```
        Validated bytes:   0
        Decrypted bytes:   0
    Secure Association received
        Accepted packets:  0
        Validated bytes:   0
        Decrypted bytes:   0
```

show security macsec statistics detail (MX480 router with MPC7E-10G)

```
user@host> show security macsec statistics xe-4/0/18 detail
Interface name: xe-4/0/18
    Secure Channel transmitted
        Encrypted packets: 10
        Encrypted bytes:   840
        Protected packets: 0
        Protected bytes:   0
    Secure Association transmitted
        Encrypted packets: 10
        Protected packets: 0
    Secure Channel received
        Accepted packets:  0
        Validated bytes:   0
        Decrypted bytes:   0
    Secure Association received
        Accepted packets:  0
        Validated bytes:   0
        Decrypted bytes:   0
    Error and debug
    Secure Channel transmitted packets
        Untagged: 0, Too long: 0
    Secure Channel received packets
        Control: 0, Tagged miss: 0
        Untagged hit: 0, Untagged: 0
        No tag: 8590007894, Bad tag: 0
        Unknown SCI: 0, No SCI: 0
        Control pass: 0, Control drop: 0
        Uncontrol pass: 0, Uncontrol drop: 0
        Hit dropped: 0, Invalid accept: 0
        Late drop: 0, Delayed accept: 0
        Unchecked: 0, Not valid drop: 0
        Not using SA drop: 0, Unused SA accept: 0
```

show security mka sessions (MX Series)

Syntax

show security mka sessions
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers. Support for MPC7E-10G introduced in Junos OS Release 16.1R1 for MX240, MX480, and MX960 routers.

Description

Display MACsec Key Agreement (MKA) session information for all interfaces. The MKA protocol is responsible for maintaining MACsec on the link, and decides which router on the point-to-point link becomes the key server.

Options

- **interface** *interface-name*—(Optional) Display the MKA session information for the specified interface only.
- none—Display the MKA session information for all interfaces.

Required Privilege Level

view

List of Sample Output

show security mka sessions on page 2172
show security mka sessions (MX480 with MPC7E-10G) on page 2172
show security mka sessions (MX480 with MPC7E-10G) on page 2173

Output Fields

Table 181 on page 2171 lists the output fields for the **show security mka sessions** command. Output fields are listed in the approximate order in which they appear.

Table 181: show security mka sessions Output Fields

Field Name	Field Description
Interface name	Name of the interface.
Member identifier	Name of the member identifier.
CAK name	Name of the connectivity association key (CAK). The CAK is configured using the **cak** keyword when configuring the pre-shared key.
Transmit interval	The transmit interval. Both ends of the point-to-point link should be configured to the same value. Default value is 2000 seconds. Possible values: 2000 through 10000 milliseconds.
Outbound SCI	Name of the outbound secure channel identifier.
Message number	Number of the last data message.
Key number	Key number.

Table 181: show security mka sessions Output Fields (continued)

Field Name	Field Description
Key server	Key server status. The router is the key server when this output is **yes**. The router is not the key server when this output is **no**.
Key server priority	Displays the priority of the key server. Lower value indicates higher priority. Use the **key-server-priority** statement to set the priority. Possible values: 0 through 255.
Latest SAK AN	Name of the latest secure association key (SAK) association number.
Latest SAK KI	Name of the latest secure association key (SAK) key identifier.
Fields for Peer list	
Member identifier	Name of the member identifier.
Hold time	Hold time, in seconds.
Message number	Number of the last data message
SCI	Name of the secure channel identifier.
Lowest acceptable PN	Number of the lowest acceptable packet number (PN).

Sample Output

show security mka sessions

```
user@host> show security mka sessions

  Interface name: xe-0/1/0
      Member identifier: 0CCBEE42F8778300F8D0C1DC
      CAK name: 1234567890
      Transmit interval: 2000(ms)
      Outbound SCI: 2C:6B:F5:9D:4B:1B/1
      Message number: 1526465    Key number: 0
      Key server: no             Key server priority: 15
      Latest SAK AN: 0           Latest SAK KI: 4F18CE25228178FD15976E4C/1
      Previous SAK AN: 0         Previous SAK KI: 000000000000000000000000/0
      Peer list
       1. Member identifier: 4F18CE25228178FD15976E4C (live)
          Message number: 1526484 Hold time: 14500 (ms)
          SCI: 2C:6B:F5:9D:3A:1B/1
          Lowest acceptable PN: 121198
```

show security mka sessions (MX480 with MPC7E-10G)

```
user@host> show security mka sessions
Interface name: xe-4/0/18
      Member identifier: FA606FD4A4C2172F0C9D9C1F
      CAK name: ABCDEF
      Transmit interval: 2000(ms)
```



```
Outbound SCI: 54:1E:56:B4:0D:3A/1
Message number: 72455      Key number: 0
Key server: no             Key server priority: 16
Latest SAK AN: 1           Latest SAK KI: 88EC3950C7D598623A406AC8/2
Previous SAK AN: 0         Previous SAK KI: 000000000000000000000000/0
Peer list
 1. Member identifier: 88EC3950C7D598623A406AC8 (live)
    Message number: 72552 Hold time: 4500 (ms)
    SCI: 54:1E:56:B3:CA:A7/1
    Lowest acceptable PN: 0
```

show security mka sessions (MX480 with MPC7E-10G)

```
user@host> show security mka sessions interface xe-1/0/7
Member identifier: 653D8911B42DAE946993B40F
      CAK name: 1111
      Transmit interval: 2000(ms)
      Outbound SCI: 54:1E:56:B3:CA:9C/1
      Message number: 179139     Key number: 0
      Key server: no             Key server priority: 16
      Latest SAK AN: 0           Latest SAK KI: 64EF352178BD1833600338F9/1
      Previous SAK AN: 0         Previous SAK KI: 000000000000000000000000/0
      Peer list
       1. Member identifier: 64EF352178BD1833600338F9 (live)
          Message number: 179175 Hold time: 4500 (ms)
          SCI: 54:1E:56:B4:0D:2F/1
          Lowest acceptable PN: 0
```

show security mka statistics (MX Series)

Syntax

show security mka statistics
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 15.1 for MX240, MX480, and MX960 routers. Support for MPC7E-10G introduced in Junos OS Release 16.1R1 for MX240, MX480, and MX960 routers.

Description

Display MACsec Key Agreement (MKA) protocol statistics.

The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see *show security macsec statistics*.

Options

- **interface** ***interface-name***—(Optional) Display the MKA information for the specified interface only.
- none—Display the MKA information for all interfaces.

Required Privilege Level

view

List of Sample Output

show security mka statistics on page 2175
show security mka statistics (MX480 routers with MPC7E-10G) on page 2175
show security mka statistics (MX480 routers with MPC7E-10G) on page 2176

Output Fields

Table 182 on page 2174 lists the output fields for the **show security mka statistics** command. Output fields are listed in the approximate order in which they appear.

Table 182: show security mka statistics Output Fields

Field Name	Field Description
Received packets	Number of received MKA control packets. This counter increments for received MKA control packets only. This counter does not increment when data packets are received.
Transmitted packets	Number of transmitted MKA packets This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.
Version mismatch packets	Number of version mismatch packets.
CAK mismatch packets	Number of Connectivity Association Key (CAK) mismatch packets. This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.

Table 182: show security mka statistics Output Fields (continued)

Field Name	Field Description
ICV mismatch packets	Number of ICV mismatched packets. This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link.
Duplicate message identifier packets	Number of duplicate message identifier packets.
Duplicate message number packets	Number of duplicate message number packets.
Duplicate address packets	Number of duplicate source MAC address packets.
Invalid destination address packets	Number of invalid destination MAC address packets.
Formatting error packets	Number of formatting error packets.
Old Replayed message number packets	Number of old replayed message number packets.

Sample Output

show security mka statistics

```
user@host> show security mka statistics

        Received packets:                      1525844
        Transmitted packets:                   1525841
        Version mismatch packets:              0
        CAK mismatch packets:                  0
        ICV mismatch packets:                  0
        Duplicate message identifier packets:  0
        Duplicate message number packets:      0
        Duplicate address packets:             0
        Invalid destination address packets:   0
        Formatting error packets:              0
        Old Replayed message number packets:   0
```

show security mka statistics (MX480 routers with MPC7E-10G)

```
user@host> show security mka statistics
 Interface name: xe-4/0/18
        Received packets:                      73009
        Transmitted packets:                   73011
        Version mismatch packets:              0
        CAK mismatch packets:                  1
        ICV mismatch packets:                  0
        Duplicate message identifier packets:  0
        Duplicate message number packets:      0
        Duplicate address packets:             0
        Invalid destination address packets:   0
        Formatting error packets:              0
        Old Replayed message number packets:   0
```


show security mka statistics (MX480 routers with MPC7E-10G)

```
user@host> show security mka statistics interface xe-1/0/7
Received packets:                    179211
        Transmitted packets:                 179186
        Version mismatch packets:            0
        CAK mismatch packets:                0
        ICV mismatch packets:                0
        Duplicate message identifier packets: 0
        Duplicate message number packets:    0
        Duplicate address packets:           0
        Invalid destination address packets:  0
        Formatting error packets:            0
        Old Replayed message number packets:  0
```

show vrrp

Syntax

show vrrp
<brief | detail | extensive | summary>
<interface *interface-name* <group *number*>>
<logical-system *logical-system-name* >
<nsr>

Release Information

Command introduced before Junos OS Release 7.4.
nsr option added in Junos OS Release 13.2.

Description

Display status information about Virtual Router Redundancy Protocol (VRRP) groups.

Options

none—(Same as **brief**) Display brief status information about all VRRP interfaces.

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

interface *interface-name* <group *number*>—(Optional) Display information and status about the specified VRRP interface and, optionally, the group number.

logical-system *logical-system-name*—(Optional) Perform this operation on a particular logical system.

nsr—(Optional) Display state replication information when graceful Routing Engine switchover (GRES) with nonstop active routing (NSR) is configured. Use only on the backup Routing Engine.

Required Privilege Level

view

Related Documentation

- *show vrrp track*
- *clear vrrp*

List of Sample Output

show vrrp on page 2183
show vrrp brief on page 2183
show vrrp detail (IPv6) on page 2183
show vrrp detail (Route Track) on page 2184
show vrrp detail (Route Track) on page 2184
show vrrp extensive on page 2184
show vrrp interface on page 2185
show vrrp nsr on page 2186
show vrrp summary on page 2187

Output Fields

Table 183 on page 2178 lists the output fields for the **show vrrp** command. Output fields are listed in the approximate order in which they appear

Table 183: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	**brief extensive** none **summary**
Interface index	Physical interface index number, which reflects its initialization sequence.	**extensive**
Groups	Total number of VRRP groups configured on the interface.	**extensive**
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	**extensive**
Interface VRRP PDU statistics	Non-errored statistics for the logical interface: • **Advertisement sent**—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • **Advertisement received**—Number of VRRP advertisement PDUs received by the interface. • **Packets received**—Number of VRRP packets received for VRRP groups on the interface. • **No group match received**—Number of VRRP packets received for VRRP groups that do not exist on the interface.	**extensive**
Interface VRRP PDU error statistics	Errored statistics for the logical interface: • **Invalid IPAH next type received**—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • **Invalid VRRP ttl value received**—Number of packets received whose IP time- to-live (TTL) value is not 255. • **Invalid VRRP version received**—Number of packets received whose VRRP version is not 2. • **Invalid VRRP pdu type received**—Number of packets received whose VRRP PDU type is not 1. • **Invalid VRRP authentication type received**—Number of packets received whose VRRP authentication is not none, simple, or md5. • **Invalid VRRP IP count received**—Number of packets received whose VRRP IP count exceeds 8. • **Invalid VRRP checksum received**—Number of packets received whose VRRP checksum does not match the calculated one.	**extensive**
Physical interface	Name of the physical interface.	**detail extensive**
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	**brief detail extensive** none
Index	Physical interface index number, which reflects its initialization sequence.	**detail extensive**
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive**

Table 183: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
VRRP-Traps	Status of VRRP traps: **Enabled** or **Disabled**.	**detail extensive**
VRRP-Version	VRRP version: **2** or **3**.	**detail extensive**
Type and Address	Identifier for the address and the address itself: • **lcl**—Configured local interface address. • **mas**—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. • **vip**—Configured virtual IP addresses.	**brief** none **summary**
Interface state/Int state/State	State of the physical interface: • **down**—The device is present and the link is unavailable. • **not present**—The interface is configured, but no physical device is present. • **unknown**—The VRRP process has not had time to query the kernel about the state of the interface. • **up**—The device is present and the link is established.	**brief extensive** none **summary**
Group	VRRP group number.	**brief extensive** none **summary**
State	The state of the interface on which VRRP is running: • **backup**—The interface is acting as the backup router interface. • **bringup**—VRRP is just starting and the physical device is not yet present. • **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • **init**—VRRP is initializing. • **master**—The interface is acting as the master router interface. • **master(ISSU)**—The master router interface is going through a unified in-service software upgrade. • **transition**—The interface is changing between being the backup and being the master router.	**extensive**
VRRP Mode	If the interface inherits its state and configuration from the active VRRP group, or if it is part of the active VRRP group. • Active—Part of the active VRRP group • Inherit—Inherits state and configuration from the active VRRP group.	**detail extensive**
Priority	Configured VRRP priority for the interface.	**detail extensive**
Advertisement interval	Configured VRRP advertisement interval.	**detail extensive**
Authentication type	Configured VRRP authentication type: **none**, **simple**, or **md5**.	**detail extensive**

Table 183: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Advertisement Threshold	A value from 1 through 15, used for setting the time when a peer should be considered down. • The time a peer is considered down is equal to the advertisement-threshold multiplied by the advertisement-interval. • (advertisement-threshold *advertisement-interval) = Peer down.	**detail extensive**
Computed Send Rate	How many protocol data units (PDUs) are generated per second. Based on the number of instances and the advertisement interval.	**detail extensive**
Preempt	Whether preemption is allowed on the interface: **yes** or **no**.	**detail extensive**
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: **yes** or **no**.	**detail extensive**
VIP count	Number of virtual IP addresses that have been configured on the interface.	**detail extensive**
VIP	List of virtual IP addresses configured on the interface.	**detail extensive**
Advertisement timer	How long, in seconds, until the advertisement timer expires.	**detail extensive**
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is **N/A**.	**detail extensive**
Virtual router uptime	How long, in seconds, that the virtual router has been up.	**detail extensive**
Master router uptime	How long, in seconds, that the master route has been up.	**detail extensive**
Virtual MAC	MAC address associated with the virtual IP address.	**detail extensive**
Tracking	Whether tracking is **enabled** or **disabled**.	**detail extensive**
Current priority	Current operational priority for being the VRRP master.	**detail extensive**
Configured priority	Configured base priority for being the VRRP master.	**detail extensive**
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. **Disabled** indicates no minimum interval.	**detail extensive**
Remaining-time	(**track** option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	**detail extensive**
Interface/Tracked interface/Track Int	Name of the tracked interface.	**detail extensive**

Table 183: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Int state/Interface state/State	Current operational state of the tracked interface: **up** or **down**.	**detail extensive**
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	**detail extensive**
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	**detail extensive**
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of **down** means that the corresponding priority cost is incurred when the interface is down.	**detail extensive**
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	**detail extensive**
Route count	The number of routes being tracked.	**detail extensive**
Route	The IP address of the route being tracked.	**detail extensive**
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	**detail extensive**
Route state	The state of the route being tracked: **up**, **down**, or **unknown**.	**detail extensive**
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	**detail extensive**
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	**detail extensive**
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	**extensive**

Table 183: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
Group VRRP PDU error statistics	Errored statistics for the VRRP group: • **Bad authentication type received**—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be **none**, **simple**, or **md5** and must be the same for all routers in the VRRP group. • **Bad password received**—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group • **Bad MD5 digest received**—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. • **Bad advertisement timer received**—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. • **Bad VIP count received**—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • **Bad VIPADDR received**—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.	**extensive**
Group state transition statistics	State transition statistics for the VRRP group: • **Idle to master transitions**—Number of times that the VRRP instance transitioned from the idle state to the master state. • **Idle to backup transitions**—Number of times that the VRRP instance transitioned from the idle state to the backup state. • **Backup to master transitions**—Number of times that the VRRP instance transitioned from the backup state to the master state. • **Master to backup transitions**—Number of times that the VRRP instance transitioned from the master state to the backup state.	**extensive**
VR state	The state of the VRRP: • **backup**—The interface is acting as the backup router interface. • **bringup**—VRRP is just starting, and the physical device is not yet present. • **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • **init**—VRRP is initializing. • **master**—The interface is acting as the master router interface. • **transition**—The interface is changing between being the backup and being the master router. NOTE: When show vrrp nsr is used on the backup Routing Engine, it displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the master Routing Engine.	**brief** none **summary**

Table 183: show vrrp Output Fields (continued)

Field Name	Field Description	Level of Output
NSR	VRRP nonstop active routing is enabled for the configured VRRP group: **yes** or **no**. NOTE: A **yes** value means that the new master Routing Engine will immediately start with the VRRP State value from the original master Routing Engine. A **no** value means that the VRRP session will: • Start afresh. • Go through asilent startup period. • Move to a backup state. • Wait for the D Timer to run out before becoming the master (only if the master has not been configured already).	**brief** none
RPD-NSR	The routing options have been set to nonstop active routing: **yes** or **no**.	**brief** none
Timer	VRRP timer information: • **A**—How long, in seconds, until the advertisement timer expires. • **D**—How long, in seconds, until the Master is Down timer expires.	**brief** none

Sample Output

show vrrp

```
user@host> show vrrp
Interface          State        Group     VR state      Timer   Type  Address
fe-0/0/0.121      up               1     master       A 1.052  lcl   fec0::12:1:1:1

                                                              vip   fe80::12:1:1:99

                                                              vip   fec0::12:1:1:99
fe-0/0/2.131      up                1    master       A 0.364  lcl   fec0::13:1:1:1

                                                              vip   fe80::13:1:1:99

                                                              vip   fec0::13:1:1:99
```

show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see show vrrp on page 2183.

show vrrp detail (IPv6)

```
user@host> show vrrp detail
Physical interface: fe-0/0/0, Unit: 121, Vlan-id: 212, Address: fec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master, VRRP Mode: Active
```



```
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::12:1:1:99,
fec0::12:1:1:99
  Advertisement timer: 1.121s, Master router: fe80::12:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled

Physical interface: fe-0/0/2, Unit: 131, Vlan-id: 213, Address: fec0::13:1:1:1/120

  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::13:1:1:99,
fec0::13:1:1:99
  Advertisement timer: 0.327s, Master router: fe80::13:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
```

show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
  Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
  Advertisement Timer: 0.469s, Master router: 101.1.1.1
  Virtual router uptime: 00:02:10, Master router uptime: 00:02:05
  Virtual Mac: 00:00:5e:00:01:01
  Tracking: disabled
```

show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-1/2/0, Unit: 0, Address: 30.30.30.30/24
  Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 100, State: master
  Priority: 150, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
  Advertisement timer: 1.218s, Master router: 30.30.30.30
  Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
  Virtual MAC: 00:00:5e:00:01:64
  Tracking: enabled
    Current priority: 150, Configured priority: 150
    Priority hold-time: disabled
    Interface tracking: disabled
    Route tracking: enabled, Route count: 1
      Route              VRF name       Route state    Priority cost
      192.168.40.0/22    default        up             30
```

show vrrp extensive

```
user@host> show vrrp extensive
Interface: ge-2/0/0.0, Interface index :65539, Groups: 1, Active :1
  Interface VRRP PDU statistics
    Advertisement sent                       :0
```



```
    Advertisement received              :0
    Packets received                    :0
    No group match received             :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received     :0
    Invalid VRRP TTL value received     :0
    Invalid VRRP version received       :0
    Invalid VRRP PDU type received      :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received      :0
    Invalid VRRP checksum received      :0

Physical interface: ge-2/0/0, Unit: 0, Address: 10.10.10.1/24
  Index: 65539, SNMP ifIndex: 648, VRRP-Traps: enabled, VRRP-Version: 3
  Interface state: up, Group: 1, State: backup, VRRP Mode: Active
  Priority: 100, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 10.10.10.2
  Dead timer: 3.078s, Master priority: 0, Master router: 10.10.10.1
  Virtual router uptime: 00:00:04
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent                  :0
    Advertisement received              :0
  Group VRRP PDU error statistics
    Bad authentication Type received    :0
    Bad password received               :0
    Bad MD5 digest received             :0
    Bad advertisement timer received    :0
    Bad VIP count received              :0
    Bad VIPADDR received                :0
  Group state transition statistics
    Idle to master transitions          :0
    Idle to backup transitions          :1
    Backup to master transitions        :0
    Master to backup transitions        :0
```

show vrrp interface

```
user@host> show vrrp interface ge-0/0/0.1
Interface: ge-0/0/0.1, Interface index :324, Groups: 2, Active :2
  Interface VRRP PDU statistics
    Advertisement sent                  :39
    Advertisement received              :0
    Packets received                    :0
    No group match received             :0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received     :0
    Invalid VRRP TTL value received     :0
    Invalid VRRP version received       :0
    Invalid VRRP PDU type received      :0
    Invalid VRRP authentication type received:0
    Invalid VRRP IP count received      :0
    Invalid VRRP checksum received      :0

Physical interface: ge-0/0/0, Unit: 1, Vlan-id: 1, Address: 101.1.1.1/24
  Index: 324, SNMP ifIndex: 623, VRRP-Traps: enabled, VRRP-Version: 2
  Interface state: up, Group: 1, State: master(ISSU), VRRP Mode: Active
  Advertisement threshold: 3, Computed send rate: 0
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Advertisement threshold: 3, Computed send rate: 0
```

```
Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 101.1.1.3
  Advertisement Timer: 0.619s, Master router: 101.1.1.1
  Virtual router uptime: 00:00:22, Master router uptime: 00:00:17
  Virtual Mac: 00:00:5e:00:01:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent                      :20
    Advertisement received                  :0
  Group VRRP PDU error statistics
    Bad authentication Type received        :0
    Bad password received                   :0
    Bad MD5 digest received                 :0
    Bad advertisement timer received        :0
    Bad VIP count received                  :0
    Bad VIPADDR received                    :0
  Group state transition statistics
    Idle to master transitions              :0
    Idle to backup transitions              :1
    Backup to master transitions            :1
    Master to backup transitions            :0
Interface: fe-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          205
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0
```

show vrrp nsr

This command is similar to **show vrrp**. Here, the **VR state** column displays the current VRRP state on the master Routing Engine, which is the future VRRP state for the backup Routing Engine. Do not use on the master Routing Engine.

NSR is yes if VRRP nonstop active routing is enabled for the configured VRRP group.

RPD-NSR is yes if the routing options have been set to nonstop active routing.

```
user@host>show vrrp nsr
Interface   State  Group  VR state  VR Mode  Type  NSR  RPD-NSR  Address
ge-1/0/1.0  up     1       master    Active   lcl  yes    yes    10.0.0.1

                                              vip                10.0.0.3

ge-1/0/1.0  up     2       master    Active   lcl  yes    yes    20.0.0.1

                                              vip                20.0.0.3

ge-1/0/1.0  up     3       master    Active   lcl  yes    yes    30.0.0.1

                                              vip                30.0.0.3

ge-1/0/1.0  up     4       master    Active   lcl  yes    yes    40.0.0.1
```

```
                                                     vip              40.0.0.3
ge-1/0/1.0   up     5       master     Active        lcl   yes     yes     50.0.0.1
                                                     vip              50.0.0.3
ge-1/0/1.0   up     1       master     Active        lcl   yes     yes     1000::1
                                                     vip
fe80::200:5eff:fe00:1
                                                     vip              1000::3
ge-1/0/1.0   up     2       master     Active        lcl   yes     yes     2000::1
                                                     vip
fe80::200:5eff:fe00:2
                                                     vip              2000::3
ge-1/0/1.0   up     3       master     Active        lcl   yes     yes     3000::1
                                                     vip
fe80::200:5eff:fe00:3
                                                     vip              3000::3
ge-1/0/1.0   up     4       master     Active        lcl   yes     yes     4000::1
                                                     vip
fe80::200:5eff:fe00:4
                                                     vip              4000::3
ge-1/0/1.0   up     5       master     Active        lcl   yes     yes     5000::1
                                                     vip
fe80::200:5eff:fe00:5
                                                     vip              5000::3
```

show vrrp summary

```
user@host> show vrrp summary
Interface       State       Group   VR state       Type    Address
ge-4/2/0.0      up              1   backup         lcl     10.57.0.2
                                                   vip     10.57.0.100
```

traceroute ethernet

Syntax

traceroute ethernet
local-mep *mep-id*
maintenance-association *ma-name*
maintenance-domain *md-name*
<ttl *value*>
<wait *seconds*>
mac-address* | *mep-id
<detail>

Release Information

Command introduced in Junos OS Release 9.0.
mep-id option introduced in Junos OS Release 9.1.
local-mep option introduced in Junos OS Release 15.1

Description

Triggers the linktrace protocol to trace the route between two maintenance points. The result of the traceroute protocol is stored in the path database. To display the path database, use the **show oam ethernet connectivity-fault-management path-database** command.

Before using the traceroute command, you can verify the remote MEP's MAC address using the **show oam ethernet connectivity-fault-management path-database** command.

Options

local-mep *mep-id*—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

detail—(Optional) Provide detailed information of the responder hostname, ingress port name, egress port name, TTL, and relay action.

mac-address—Destination unicast MAC address of the remote maintenance point.

mep-id—MEP identifier of the remote maintenance point. The range of values is 1 through 8191.

maintenance-association *ma-name*—Specifies an existing maintenance association from the set of configured maintenance associations.

maintenance-domain *md-name*—Specifies an existing maintenance domain from the set of configured maintenance domains.

ttl *value*—Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.

Required Privilege Level

network

List of Sample Output

traceroute ethernet on page 2190

traceroute ethernet detail on page 2190

Output Fields Table 184 on page 2189 lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 184: traceroute ethernet Output Fields

Field Name	Field Description
Linktrace to	MAC address of the destination maintenance point.
Interface	Local interface used to send the linktrace message (LTM).
Maintenance Domain	Maintenance domain specified in the traceroute command.
Level	Maintenance domain level configured.
Maintenance Association	Maintenance association specified in the traceroute command.
Local Mep	The local maintenance end point identifier.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all Maintenance Domains. Use the transaction identifier to match an incoming linktrace response (LTR), with a previously sent LTM.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag node responding to the LTM or the source MAC address of the LTR.
Next-hop MAC address	MAC address of the egress interface of the node to which the LTM is forwarded or the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.
Responder Hostname	The hostname of the responding router. A valid hostname is received only when the responding system is a Juniper Networks router.
Ingress port name	The port name for ingress connections.
Egress port name	The port name for egress connections.

Table 184: traceroute ethernet Output Fields (continued)

Field Name	Field Description
Flags	The configurable flags can include: • H— Hardware only, incoming LT frame has hardware bit set. • T— Terminal MEP, responder is a terminating MEP. • F— FWD yes, LTM frame is relayed further.
Relay Action	The associated relay action. Relay action can be one of the following: • RlyHit— Relay hit; target MAC address matches the MP mac address. • RlyFDB— Relay FDB; output port decided by consulting forwarding database. • RlyMPDB— Relay MIP; output port decided by consulting MIP database.

Sample Output

traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:01:02:03:04:05
 Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

       Hop     TTL       Source MAC address        Next hop MAC address
       Transaction Identifier:100001
       1       63        00:00:aa:aa:aa:aa         00:00:ab:ab:ab:ab
       2       62        00:00:bb:bb:bb:bb         00:00:bc:bc:bc:bc
       3       61        00:00:cc:cc:cc:cc         00:00:cd:cd:cd:cd
       4       60        00:01:02:03:04:05         00:00:00:00:00:00
```

traceroute ethernet detail

```
user@host> run traceroute ethernet maintenance-domain md6 maintenance-association ma6
mep 101 detail
Linktrace to 00:00:5E:00:53:CC, Interface : ge-1/0/0.1
Maintenance Domain: md6, Level: 6
Maintenance Association: ma6, Local Mep: 201
Transaction Identifier: 2077547465

Legend for RelayAction:
RlyHit  -- Relay hit, Target MAC address matches the MP mac address
RlyFDB  -- Relay FDB, output port decided by consulting FDB database
RlyMPDB -- Relay MIP, output port decided by consulting MIP database

Legend for Flags:
H -- Hardware only,incoming LT frame has hardware bit set
T -- Terminal MEP, responder is a terminating MEP
F -- FWD yes, LTM frame is relayed further

TTL   Responder Hostname    Ingress port name     Egress port name
RelayAction
```

```
Responder Service      Ingress MAC address   Egress MAC address     Flags
62    host1                  ge-1/0/0.1            ge-2/3/0.1              RlyFDB
br1                    00:00:5E:00:53:00     00:00:5E:00:53:A0      HF-
63    host2                 ge-2/3/0.1            ge-1/0/0.1              RlyFDB
br1                    00:00:5E:00:53:AA     00:00:5E:00:53:A2      HF-
61    host3                 ge-1/0/0.1            --:--                   RlyHit
br1                    00:00:5E:00:53:B0          --:--             H-T
```

