Junos® OS

Ethernet Interfaces Feature Guide for Routing Devices

Release

14.1

Modified: 2015-11-09

Copyright © 2015, Juniper Networks, Inc.

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

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

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

Junos® OS Ethernet Interfaces Feature Guide for Routing Devices
14.1
Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Part 6 Index

Copyright © 2015, Juniper Networks, Inc.

List of Figures

 Copyright © 2015, Juniper Networks, Inc.

List of Tables

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

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 http://www.juniper.net/techpubs/.

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

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

Supported Platforms

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

- ACX Series
- M Series
- MX Series
- T Series
- J 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.

Copyright © 2015, Juniper Networks, Inc.

```
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 the *CLI User Guide*.

Documentation Conventions

Table 1 on page xxxix 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 xl defines the text and syntax conventions used in this guide.

Copyright © 2015, Juniper Networks, Inc.

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.	

GUI Conventions

Copyright © 2015, Juniper Networks, Inc.

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

Copyright © 2015, Juniper Networks, Inc.

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

 Copyright © 2015, Juniper Networks, Inc.

PART 1

Ethernet Interfaces

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

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

Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Configuring Ethernet Physical Interface Properties on page 8
- Configuring Gigabit Ethernet Interfaces on J Series Services Routers on page 6
- MX Series Router Interface Identifiers on page 4
- Enabling Ethernet MAC Address Filtering on page 15
- Configuring Ethernet Loopback Capability on page 19
- Configuring Flow Control on page 13
- Ignoring Layer 3 Incomplete Errors on page 20
- Configuring the Link Characteristics on Ethernet Interfaces on page 14
- Configuring Gratuitous ARP on page 20
- Adjusting the ARP Aging Timer on page 21
- Configuring the Interface Speed on Ethernet Interfaces on page 12
- Configuring the Ingress Rate Limit on page 13
- Configuring Multicast Statistics Collection on Ethernet Interfaces on page 22
- Configuring Weighted Random Early Detection on page 22
- *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*

 Copyright © 2015, Juniper Networks, Inc.

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. J Series Gigabit Ethernet interfaces do not support 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 Gigabit Ethernet Interfaces on J Series Services Routers

The J Series routers with multiport Gigabit Ethernet uPIMs supports Ethernet access switching. This functionality provides the ability to switch traffic at Layer 2 in addition to routing traffic at Layer 3.

J Series routers with multiport Gigabit Ethernet uPIMs can be deployed in branch offices as an access or desktop switch with integrated routing capability. The multiport Gigabit Ethernet uPIM provides Ethernet switching, while the Routing Engine provides routing functionality.

Copyright © 2015, Juniper Networks, Inc.

Routed traffic is forwarded from any port of the multiport Gigabit Ethernet uPIM to the WAN interface. Switched traffic is forwarded from one port of the multiport Gigabit Ethernet uPIM to another port on the same the multiport Gigabit Ethernet uPIM. Switched traffic is not forwarded from a port on one multiport Gigabit Ethernet uPIM to a port on a different multiport Gigabit Ethernet uPIM. For more information about configuring the multiport Gigabit Ethernet uPIM switching mode, see the *Junos OS Administration Library for Routing Devices*.

In access switching mode, only one physical interface is configured for the entire multiport Gigabit Ethernet uPIM. The single physical interface serves as a Virtual Router Interface (VRI). Configuration of the physical port characteristics is done under the single physical interface.

To configure multiport Gigabit Ethernet uPIM Ethernet port properties, include the **switch-port** statement at the **[edit interfaces ge-*pim*/0/0]** hierarchy level:

```
[edit interfaces ge-pim/0/0]
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed 1g;
    link-mode (full-duplex | half-duplex);
  }
}
```

Access switching mode is supported on the 6-port, 8-port, and 16-port Gigabit Ethernet uPIMs.

The multiport Gigabit Ethernet uPIMs are supported on the J2320, J2350, J4350, and J6350 Services Routers.

The 6-port and 8-port multiport Gigabit Ethernet uPIM occupies a single slot and can be installed in any slot. Because the 16-port Gigabit Ethernet uPIM is two slots high, you cannot install a 16-port uPIM in the top slots (slots 1 and 4). 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.

Example: Configuring J Series Services Router Switching Interfaces

Configure a single physical interface for the uPIM and set the port parameters for port 0 and port 1:

```
[edit interfaces]
ge-2/0/0 {
  switch-options {
    switch-port 0 {
      no-auto-negotiation;
      speed 1g;
      link-mode full-duplex;
    }
    switch-port 1 {
      no-auto-negotiation;
      speed 10m;
      link-mode half-duplex;
```

Copyright © 2015, Juniper Networks, Inc.

```
    }
  }
}
```

Related Documentation

- switch-options on page 864
- switch-port on page 865
- speed on page 853
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Ethernet Physical Interface Properties

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]
link-mode (full-duplex | half-duplex);
speed (10m | 100m);
vlan-tagging;
fastether-options {
  802.3ad aex (primary | backup);
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  ingress-rate-limit rate;
  (loopback | no-loopback);
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```


Copyright © 2015, Juniper Networks, Inc.

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, the J Series Gigabit Ethernet uPIM 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 J Series Gigabit Ethernet uPIM interfaces and 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.

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]
gigether-options {
  802.3ad aex (primary | backup);
  auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  no-auto-mdix;
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

Additionally, 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]
optics-options {
  wavelength nm;
```

Copyright © 2015, Juniper Networks, Inc.

```
}
```

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). For more information, see "Example: Configuring Gigabit Ethernet Interfaces" on page 6.

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
gigether-options {
  802.3ad aex (primary | backup);
  auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  (source-filtering | no-source-filtering);
  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;
        }
      }
    }
  }
  native-vlan-id number;
}
```

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. For more information, see "10-Gigabit Ethernet Framing Overview" on page 257.

 Copyright © 2015, Juniper Networks, Inc.

```
[edit interfaces]
xe-0/0/0 {
   framing {
      (lan-phy | wan-phy);
   }
}
```

To configure OAM 802.3ah 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;
            }
         }
      }
   }
}
```

To configure Gigabit Ethernet IQ-specific logical interface properties, include the **input-vlan-map**, **output-vlan-map**, **layer2-policer**, and **vlan-tags** statements:

```
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;
}
layer2-policer {
   input-policer policer-name;
   input-three-color policer-name;
   output-policer policer-name;
   output-three-color policer-name;
}
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
```

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*]**

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

```
[edit interfaces aex]
aggregated-ether-options {
  ethernet-switch-profile {
    tag-protocol-id tpid;
  }
  (flow-control | no-flow-control);
  lacp mode {
    periodic interval;
  }
  link-protection;
  link-speed speed;
  (loopback | no-loopback);
  minimum-links number;
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

Related Documentation

- Example: Configuring Gigabit Ethernet Interfaces on page 6
- 10-Gigabit Ethernet Framing Overview on page 257
- 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**), the J Series Gigabit Ethernet uPIM interfaces, 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 J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps **(10m | 100m | 1g)**. 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.

NOTE: 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.

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

Copyright © 2015, Juniper Networks, Inc.

NOTE: 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.

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

```
[edit interfaces interface-name]
speed (10m | 100m | 1g);
```

Related Documentation

- speed on page 853
- 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 708
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

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;
```

Copyright © 2015, Juniper Networks, Inc.

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 690
- Ethernet Interfaces Overview on page 3

- *EX Series Switches Interfaces Overview*
- *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. J Series Gigabit Ethernet interfaces and Fast Ethernet interfaces, except the J Series ePIM 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.

Copyright © 2015, Juniper Networks, Inc.

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

NOTE: On a J Series ePIM Fast Ethernet interface, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."

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);

Related Documentation

- link-mode on page 735
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling Ethernet MAC Address Filtering

By default, source address filtering is disabled. 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, which blocks all incoming packets to an interface.

NOTE: Source address filtering is not supported on J Series Services Routers.

To enable the filtering, include the **source-filtering** statement:

source-filtering;

Copyright © 2015, Juniper Networks, Inc.

To explicitly disable filtering, include the **no-source-filtering** statement:

```
no-source-filtering;
```

You can include these statements 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.

Filtering Specific MAC Addresses

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** statement:

```
source-address-filter {
  mac-address;
  <additional-mac-address>;
}
```

You can include these statements 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.

Copyright © 2015, Juniper Networks, Inc.

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 316.

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 56.

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 316.

Related Documentation

- **source-address-filter on page 851**
- Configuring MAC Address Filtering on page 316
- Configuring LACP for Aggregated Ethernet Interfaces on page 56
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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;
    }
  }
}
```

 Copyright © 2015, Juniper Networks, Inc.

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 (Gigabit Ethernet) on page 1309

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 746
- Ethernet Interfaces Overview on page 3
- *EX Series Switches Interfaces Overview*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 707
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests provide duplicate IP address detection. A gratuitous ARP request 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. 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.

By default, the router or switch responds to gratuitous ARP requests. On Ethernet interfaces, you can disable responses to gratuitous ARP requests. To disable responses to gratuitous ARP requests, include the **no-gratuitous-arp-request** statement at the **[edit interfaces *interface-name*]** hierarchy level:

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

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. On Ethernet interfaces, you can enable handling of gratuitous ARP replies on

 Copyright © 2015, Juniper Networks, Inc.

a specific interface by including the **gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
gratuitous-arp-reply;
```

To restore the default behavior, include the **no-gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
no-gratuitous-arp-reply;
```

Related Documentation

- gratuitous-arp-reply on page 700
- no-gratuitous-arp-request on page 776
- Ethernet Interfaces Overview on page 3
- *EX Series Switches Interfaces Overview*
- *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 for Routing Devices*.

Related Documentation

- *arp (System)*
- 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*.
- 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 104
- Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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*

Copyright © 2015, Juniper Networks, Inc.

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 18
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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;
    }
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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**).

Copyright © 2015, Juniper Networks, Inc.

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 18
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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: I2.0 based M120 routers and I3.0 based M320 routers with the PICs in the preceding list support passive monitoring starting with Junos OS Release 9.5. 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.**

Related Documentation

- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Enabling Passive Monitoring on Ethernet Interfaces

On Ethernet interfaces, enable packet flow monitoring by including the **passive-monitor-mode** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
passive-monitor-mode;
```

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.

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:

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

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]
receive-options-packets;
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*.

Copyright © 2015, Juniper Networks, Inc.

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*

Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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.

 Copyright © 2015, Juniper Networks, Inc.

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*

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 6

Configuring Aggregated Ethernet Interfaces for Increased Throughput and Link Redundancy

- Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces on page 104
- Deleting an Aggregated Ethernet Interface on page 105

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*.

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.

 Copyright © 2015, Juniper Networks, Inc.

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 3 on page 39 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 3: 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)

Copyright © 2015, Juniper Networks, Inc.

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.

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, PICs, or MPCs.

Configuration Guidelines for Aggregated Ethernet Interfaces

- 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*.
- 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 303.
- 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:

 Copyright © 2015, Juniper Networks, Inc.

- 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 3 on page 39 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.

Related Documentation

- inner-tag-protocol-id on page 709
- inner-vlan-id on page 710
- pop-pop on page 809
- pop-swap on page 810
- push-push on page 824
- swap-push on page 862
- swap-swap on page 863
- Configuring Mixed Aggregated Ethernet Links on page 45
- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 303
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

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 ae*x***

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.

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

 Copyright © 2015, Juniper Networks, Inc.

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 47.

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 50
- Deleting an Aggregated Ethernet Interface on page 105
- Aggregated Ethernet Interfaces Overview on page 38

- *Ethernet Interfaces Feature Guide for Routing Devices*

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
- 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 38
- Configuring Junos OS for Supporting Aggregated Devices on page 48

 Copyright © 2015, Juniper Networks, Inc.

Configuring Mixed Aggregated Ethernet Links

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 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.

To configure member links of mixed rate or mixed mode aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers, you need to configure the **mixed** option for the [**edit interfaces ae*x* aggregated-ether-options link-speed**] statement.

To configure mixed aggregated Ethernet interfaces:

1. Configure the number of aggregated logical devices available to the router:

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

   For example:

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count 3
   ```

2. Configure the minimum number of links that is required for the aggregated Ethernet interface to be labeled *up*:

   ```
   [edit interfaces]
   user@host# set aex aggregated-ether-options minimum-links number
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options minimum-links 2
   ```

Copyright © 2015, Juniper Networks, Inc.

NOTE: By default, only one link needs to be up for the bundle to be labeled *up*.

3. Configure the **link-speed** statement. Specify the **mixed** option for the **link-speed** statement to indicate the mixed aggregated Ethernet bundle configuration.

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

   For example:

   [edit interfaces]
   user@host# **set ae0 aggregated-ether-options link-speed mixed**

   NOTE: It is mandatory to configure the **mixed** option when configuring the mixed aggregated Ethernet bundle on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4). On other supported platforms, if the **[edit interfaces ae*x* aggregated-ether-options link-speed]** statement is not configured, the mixed configuration is applied by default.

4. Configure the members links of the aggregated Ethernet bundle:

   [edit interfaces]
   user@host# **set *interface-name* gigether-options 802.3ad ae*x***

   For example:

   [edit interfaces]
   user@host# **set xe-0/0/1 gigether-options 802.3ad ae0**
   user@host# **set et-1/1/0 gigether-options 802.3ad ae0**
   user@host# **set ce-1/1/1 gigether-options 802.3ad ae0**

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

   [edit interfaces]
   user@host# **set ae*x* unit *number* family inet address *address***

   For example:

   [edit interfaces]
   user@host# **set ae0 unit 0 family inet address 100.100.100.1/30**

6. Commit the configuration:

   [edit]
   user@host# **commit**

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 38
- Configuring Aggregated Ethernet Link Speed on page 51
- **link-speed on page 738**

Copyright © 2015, Juniper Networks, Inc.

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 14.0.100.50/24;
    }
  }
  unit 1 {
    vlan-id 1024;
    family inet {
      address 14.0.101.50/24;
    }
  }
  unit 2 {
    vlan-id 1025;
    family inet {
      address 14.0.102.50/24;
    }
  }
  unit 3 {
    vlan-id 4094;
    family inet {
      address 14.0.103.50/24;
    }
```

Copyright © 2015, Juniper Networks, Inc.

```
  }
}
```

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;
  }
}
```

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**

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.

 Copyright © 2015, Juniper Networks, Inc.

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).

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.

Copyright © 2015, Juniper Networks, Inc.

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;
```

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 42
- *Ethernet Interfaces Feature Guide for Routing Devices*
- Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers on page 43
- *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**

1. Specify that you want to access the aggregated Ethernet configuration on the device.

   user@host# **edit chassis aggregated-devices ethernet**

 Copyright © 2015, Juniper Networks, Inc.

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.

For information about E Series routers, see *Understanding Aggregated Ethernet Interfaces and LACP*.

Related Documentation

- For information about physical links, see Configuring an Aggregated Ethernet Interface on page 42
- For a sample configuration, see Example: Configuring Aggregated Ethernet Interfaces on page 47
- *Ethernet Interfaces Feature Guide for Routing Devices*
- For information about configuring aggregated devices, see the *Junos OS Administration Library for Routing Devices*.

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, 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)

Copyright © 2015, Juniper Networks, Inc.

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, and TX Matrix Plus routers, you need to configure the **mixed** option for the [**edit interfaces** ***aex*** **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.

Copyright © 2015, Juniper Networks, Inc.

Aggregated Ethernet links on T Series routers 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.

Related Documentation

- **aggregated-ether-options on page 617**
- Configuring Mixed Aggregated Ethernet Links on page 45
- *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**.

Copyright © 2015, Juniper Networks, Inc.

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**.

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 617
- minimum-links on page 764
- *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 136.

Related Documentation

- vlan-id on page 900
- vlan-tagging on page 908

Configuring Untagged Aggregated Ethernet Interfaces

When you configure an untagged Aggregated Ethernet interface, the existing rules for untagged interfaces apply. These rules are as follows:

- 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.
- You cannot include the **vlan-id** statement in the configuration of the logical interface.

Table 4 on page 55 lists untagged aggregated Ethernet and LACP support by PIC and router.

Copyright © 2015, Juniper Networks, Inc.

Table 4: 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.

For more information about configuring LACP, see “Configuring LACP for Aggregated Ethernet Interfaces” on page 56.

Example: Configuring Untagged Aggregated Ethernet Interfaces

Configure an untagged aggregated Ethernet interface by omitting the **vlan-tagging** and **vlan-id** statements from the configuration:

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  # vlan-tagging; OMIT FOR UNTAGGED AE CONFIGURATIONS
  unit 0 {
    # vlan-id 100; OMIT FOR UNTAGGED AE CONFIGURATIONS
    family inet {
      address 13.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 13.1.1.4;
          priority 200;
        }
      }
    }
```

Copyright © 2015, Juniper Networks, Inc.

```
  }
}
```

Related Documentation

- For more information about configuring LACP, see Configuring LACP for Aggregated Ethernet Interfaces on page 56.
- *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.

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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;
}
```

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 on page 58
- Configuring LACP Link Protection on page 58
- Tracing LACP Operations on page 61
- LACP Limitations on page 62
- Example: Configuring Aggregated Ethernet LACP on page 62

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: 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 "Enabling Ethernet MAC Address Filtering" on page 15.

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 for Routing Devices*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see "Configuring Untagged Aggregated Ethernet Interfaces" on page 54.

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

NOTE: When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.

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.

Copyright © 2015, Juniper Networks, Inc.

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: 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

Copyright © 2015, Juniper Networks, Inc.

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;
```

You must 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*.

 Copyright © 2015, Juniper Networks, Inc.

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: **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.**

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;
}
```

You can specify the following flags in the **protocols lacp traceoptions** statement:

- **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

For general information about tracing, see the tracing and logging information in the *Junos OS Administration Library for Routing Devices*.

LACP Limitations

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.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

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;
        }
      }
    }
```


Copyright © 2015, Juniper Networks, Inc.

```
  }
}
```

Configure aggregated Ethernet LACP over an untagged interface:

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;
        }
      }
    }
  }
}
```

Related Documentation

- lacp on page 726
- link-protection on page 737
- traceoptions on page 878
- *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.

Copyright © 2015, Juniper Networks, Inc.

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**

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**

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:

 Copyright © 2015, Juniper Networks, Inc.

user@host> **request interface revert ae*x***

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**

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 63

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 ae*X* 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*.

Copyright © 2015, Juniper Networks, Inc.

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

This section describes configuration of symmetrical load balancing on an 802.3ad link aggregation group (LAG) 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 1 on page 68.

Figure 1: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

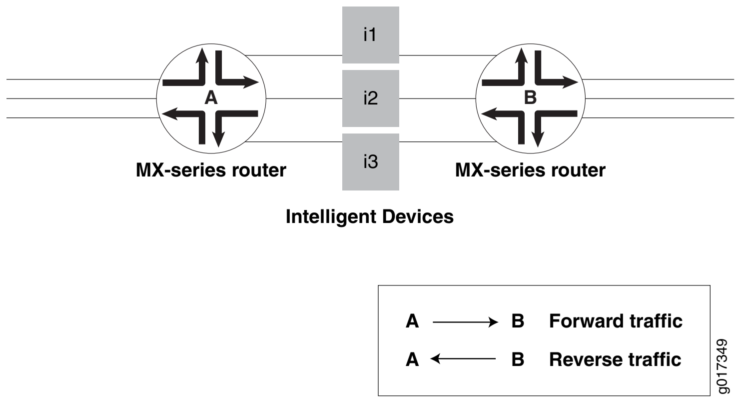

Router A is configured with symmetric hash and Router B is configured with symmetric hash complement. Thus, for a given flow *fx*, post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow *fx* is from Router B to Router A through the same i2 device as its hashing (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.

 Copyright © 2015, Juniper Networks, Inc.

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 for Routing Devices*.

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.

 Copyright © 2015, Juniper Networks, Inc.

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 2 on page 72, 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.

Copyright © 2015, Juniper Networks, Inc.

Figure 2: Traffic Polarization on Cascaded Routers When Symmetrical Load Balancing in Enabled on Trio-based MPCs

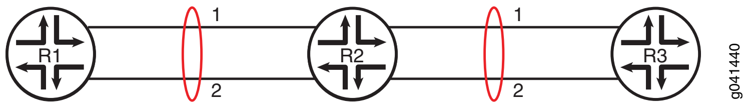

This is unlike having a random link selection algorithm, where a flow might use Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, and Link 2 of the aggregated Ethernet bundle between Device R2 and Device R3.

- 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

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;
}
```

 Copyright © 2015, Juniper Networks, Inc.

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 for Routing Devices*.

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)

Copyright © 2015, Juniper Networks, Inc.

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 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.

 Copyright © 2015, Juniper Networks, Inc.

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 78

Example: Configuring Aggregated Ethernet Load Balancing

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

Copyright © 2015, Juniper Networks, Inc.

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 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.

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

Copyright © 2015, Juniper Networks, Inc.

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 3: Aggregated Ethernet Load Balancing

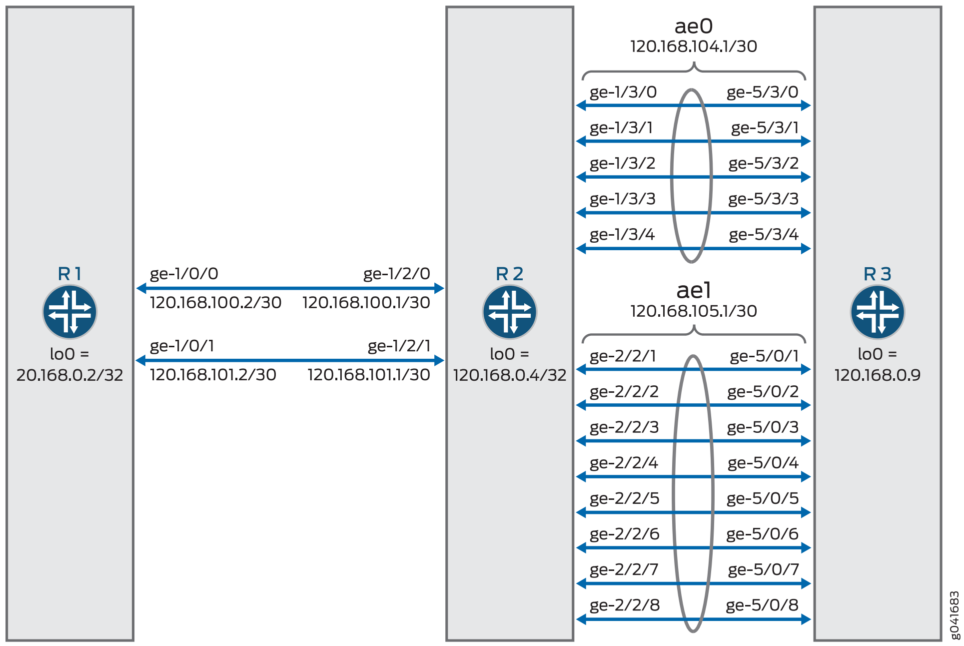

Configuration

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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.

 Copyright © 2015, Juniper Networks, Inc.

```
[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
```

Copyright © 2015, Juniper Networks, Inc.

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
        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;
```


Copyright © 2015, Juniper Networks, Inc.

```
        }
        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 89

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.

Configuring Adaptive Load Balancing

This topic describes how to configure adaptive load balancing on PTX Series Packet Transport Routers. 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 router 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 bundle on PTX Series routers:

1. Enable adaptive load balancing on the AE bundle:

   [edit interfaces ae-*x* aggregated-ether-options load-balance]
   user@router# **set adaptive**

 Copyright © 2015, Juniper Networks, Inc.

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 73
- Example: Configuring Aggregated Ethernet Load Balancing on page 78
- **adaptive on page 608**

Understanding Independent Micro BFD Sessions for LAG

Copyright © 2015, Juniper Networks, Inc.

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

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 of the member links 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: 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.
- Specify the neighbor in a BFD session. The neighbor address can be either an IPv4 or an IPv6 address.

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.

Related Documentation

- *authentication*
- **bfd-liveness-detection on page 633**
- *detection-time*
- *transmit-interval*
- *Configuring Independent Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG on page 95

Copyright © 2015, Juniper Networks, Inc.

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 MPC cards.
- T Series routers with Type 4 FPC or Type 5 FPC.
- 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 and AE1 that are assigned four and two IP addresses respectively. Configure BFD sessions on the aggregated Ethernet interfaces on both ends of the connection. This example verifies that independent micro BFD sessions are active in the output.

Topology

Figure 4 on page 95 shows the sample topology.

Figure 4: Configuring Independent Micro BFD Session for LAG

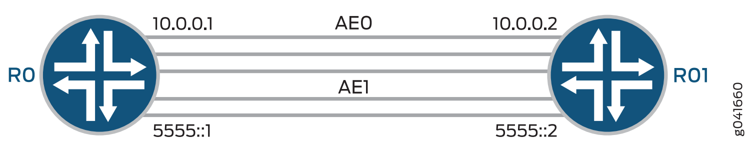

Configuration

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 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 lo0 unit 0 family inet address 10.255.106.107/32
```

```
set 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 ae0 unit 0 family inet address 10.0.0.5/30
set interfaces ae0 unit 0 family inet address 10.0.0.9/30
set interfaces ae0 unit 0 family inet address 10.0.0.13/30
set interfaces ae0 unit 0 family inet6 address 1111::1/126
set interfaces ae0 unit 0 family inet6 address 2222::1/126
set interfaces ae0 unit 0 family inet6 address 3333::1/126
set interfaces ae0 unit 0 family inet6 address 4444::1/126
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 unit 0 family inet address 50.0.0.1/30
set interfaces ae1 unit 0 family inet address 51.0.0.1/29
set interfaces ae1 unit 0 family inet6 address 5555::1/126
set interfaces ae1 unit 0 family inet6 address 6666::1/126
set routing-options nonstop-routing
set routing-options static route 30.30.30.0/30 next-hop 10.0.0.2
set routing-options static route 3ffe::1:2/126 next-hop 1111::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 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 ge-1/1/8 unit 0 family inet address 30.30.30.1/30
set interfaces ge-1/1/8 unit family inet6 address 3ffe::1:2/126
set lo0 unit 0 family inet address 10.255.106.102/32
set 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 neighbor
  10.255.106.107
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address
  10.255.106.102
set interfaces ae0 unit 0 family inet address 10.0.0.2/30
set interfaces ae0 unit 0 family inet address 10.0.0.6/30
set interfaces ae0 unit 0 family inet address 10.0.0.10/30
```

 Copyright © 2015, Juniper Networks, Inc.

```
set interfaces ae0 unit 0 family inet address 10.0.0.14/30
set interfaces ae0 unit 0 family inet6 address 1111::2/126
set interfaces ae0 unit 0 family inet6 address 2222::2/126
set interfaces ae0 unit 0 family inet6 address 3333::2/126
set interfaces ae0 unit 0 family inet6 address 4444::2/126
set interfaces ae0 aggregated-ether-options minimum-links 2
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 unit 0 family inet address 50.0.0.2/30
set interfaces ae1 unit 0 family inet address 51.0.0.2/29
set interfaces ae1 unit 0 family inet6 address 5555::2/126
set interfaces ae1 unit 0 family inet6 address 6666::2/126
set routing-options static route 20.20.20.0/30 next-hop 10.0.0.1
set routing-options static route 3ffe::1:1/126 next-hop 1111::1/126
```

Configuring 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 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 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 four IP addresses 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
user@R0# set ae0 unit 0 family inet address 10.0.0.5/30
user@R0# set ae0 unit 0 family inet address 10.0.0.9/30
user@R0# set ae0 unit 0 family inet address 10.0.0.13/30
```

```
user@R0# set ae0 unit 0 family inet6 address 1111::1/126
user@R0# set ae0 unit 0 family inet6 address 2222::1/126
user@R0# set ae0 unit 0 family inet6 address 3333::1/126
user@R0# set ae0 unit 0 family inet6 address 4444::1/126
```

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 static route 3ffe::1:2/126 next-hop 1111::2
```

```
user@R0# set static route
```

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 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 two IP addresses 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 inet address 50.0.0.1/30
user@R0# set ae1 unit 0 family inet address 51.0.0.1/29
```

```
user@R0# set ae1 unit 0 family inet6 address 5555::1/126
user@R0# set ae1 unit 0 family inet6 address 6666::1/126
```

8. Configure BFD for the aggregated Ethernet interface ae1.

Copyright © 2015, Juniper Networks, Inc.

```
[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
```

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;
         address 10.0.0.5/30;
         address 10.0.0.9/30;
         address 10.0.0.13/30;
      }
      family inet6 {
         address 1111::1/126;
         address 2222::1/126;
         address 3333::1/126;
         address 4444::1/126;
      }
   }
}
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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
        family inet {
            address 50.0.0.1/30;
            address 51.0.0.1/29;
        }
        family inet6 {
            address 5555::1/126;
            address 6666::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 1111::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 101
- Viewing Detailed BFD Events on page 103

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
                                                  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 10.0.0.2 (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: 10.0.0.2
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 10.0.0.2 is up
```

Copyright © 2015, Juniper Networks, Inc.

```
 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 633
- *detection-time*
- *Configuring Independent Micro BFD Sessions for LAG*
- Understanding Independent Micro BFD Sessions for LAG on page 92

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*

 Copyright © 2015, Juniper Networks, Inc.

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 42
- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 50
- Aggregated Ethernet Interfaces Overview on page 38
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 7

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.

Copyright © 2015, Juniper Networks, Inc.

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 5: Connections Terminating on Single PE

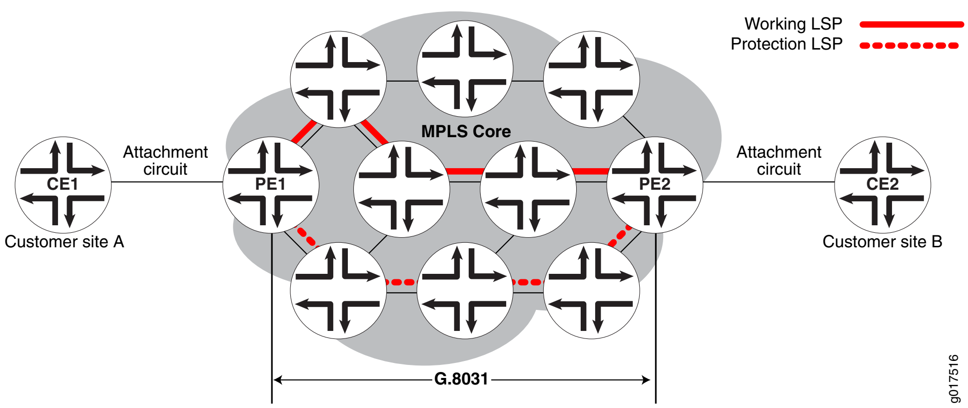

In the scenario diagramed in Figure 5 on page 108, a Virtual Private Wire Service (VPWS) is provisioned between customer sites A and B using a single pseudowire (layer 2 circuit) in the core network, and two Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) are provisioned, one for the working path and the other one for the protection

 Copyright © 2015, Juniper Networks, Inc.

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 6: Connections Terminating on a Different PE

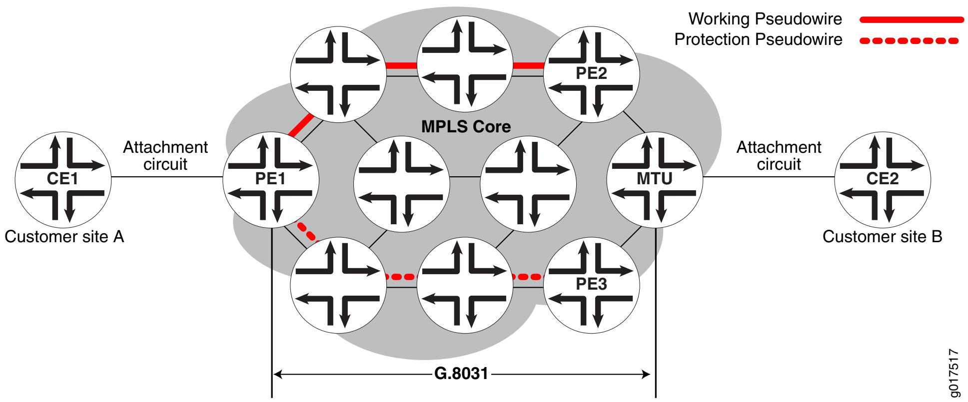

In the scenario represented in Figure 6 on page 109, a VPWS is provisioned between customer sites A and B using two pseudowires (layer 2 circuit) in the core network, one for the working path and the other for the protection path. CFM CCM will be used to monitor the status of each pseudowire.

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.

Copyright © 2015, Juniper Networks, Inc.

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 110
- Example: Configuring Protection Switching Between Psuedowires on page 111

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 7: Understanding APS Events

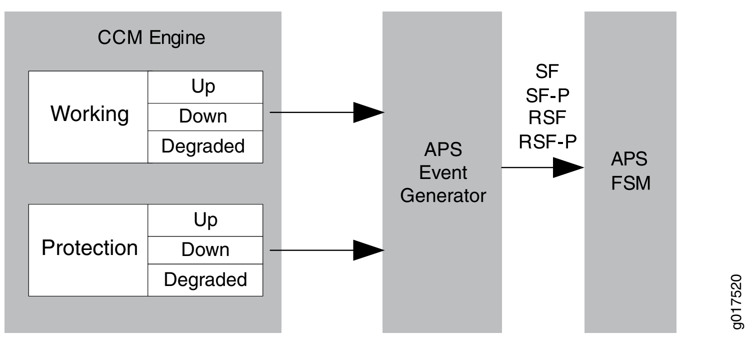

As show in Figure 7 on page 110, the APS event generator generates the following APS events based on the status of the working and protection paths:

- **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

Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 107
- Example: Configuring Protection Switching Between Psuedowires on page 111

Example: Configuring Protection Switching Between Psuedowires

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 8 on page 111.

Figure 8: Topology of a Network Using VPWS Psuedowires

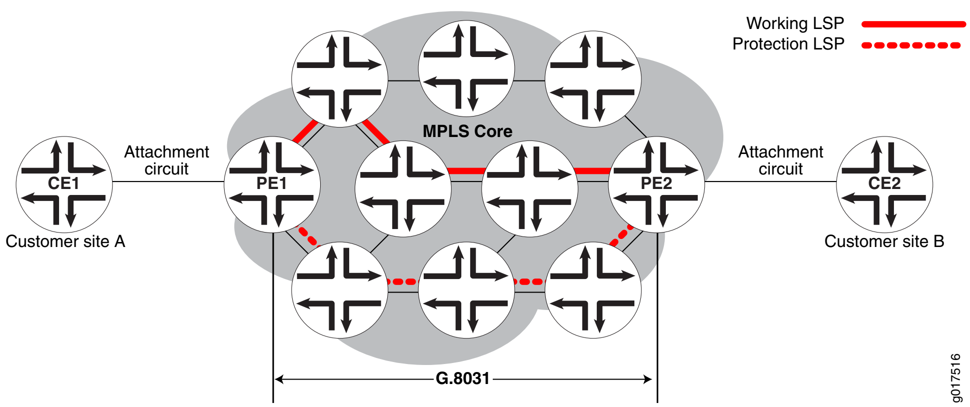

The following definitions describe the meaning of the device abbreviations used in Figure 8 on page 111.

- 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;
      }
```


Copyright © 2015, Juniper Networks, Inc.

```
  }
}
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 107
- Mapping of CCM Defects to APS Events on page 110

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 8

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.

The following standards provide detailed information on Ethernet ring protection switching:

- IEEE 802.1Q - 1998
- IEEE 802.1D - 2004
- IEEE 802.1Q - 2003
- Draft ITU-T Recommendation G.8032/Y.1344, *Ethernet Ring protection switching*
- 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*.

Copyright © 2015, Juniper Networks, Inc.

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 116
- Configuring Ethernet Ring Protection Switching on page 121
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 121
- *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

Copyright © 2015, Juniper Networks, Inc.

- WTR—Wait to restore
- 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. This node also initiates the RAPS message.

Ring Node States

There are three 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, 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.

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. 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.

Failure Detection

Ethernet ring operation depends on quick and accurate failure detection. The failure condition *signal failure (SF)* is supported. For SF detection, an Ethernet continuity check

MEP must be configured for each ring link. For fast protection switching, a 10-ms transmission period for this MEP group is supported. OAM monitors the MEP group's MA and reports SF or SF clear events to the Ethernet ring control module. For this MEP group, the action profile must be configured to update the interface device IFF_LINKDOWN flag. OAM updates the IFF_LINKDOWN flag to notify the Ethernet ring control module.

Logical Ring

This 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.

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.

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.

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 9 on page 118, 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 10 on page 118. 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 9: Protocol Packets from the Network to the Router

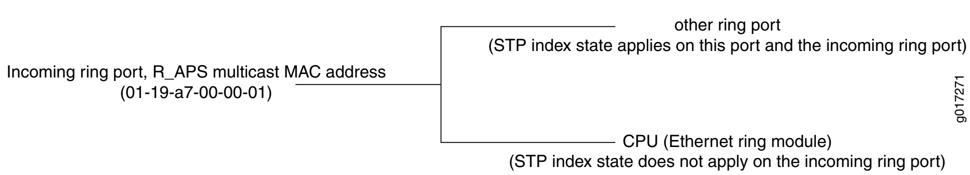

Figure 10: Protocol Packets from the Router or Switch to the Network

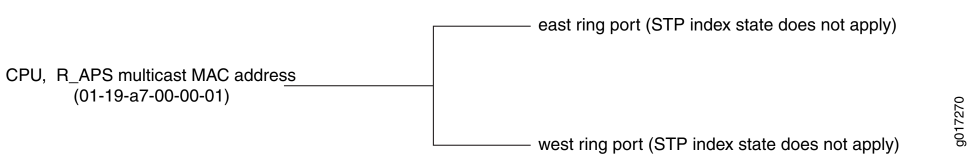

 Copyright © 2015, Juniper Networks, Inc.

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

Copyright © 2015, Juniper Networks, Inc.

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.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). However, interconnection of multiple rings is not supported in this release. The interconnection of two rings means that two rings may share the same link or share the same node.

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 such as STP. 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.

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.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 115
- Configuring Ethernet Ring Protection Switching on page 121
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 121

- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 115
- Understanding Ethernet Ring Protection Switching Functionality on page 116
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 121

- *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 11 on page 122.

Figure 11: Example of a Three-Node Ring Topology

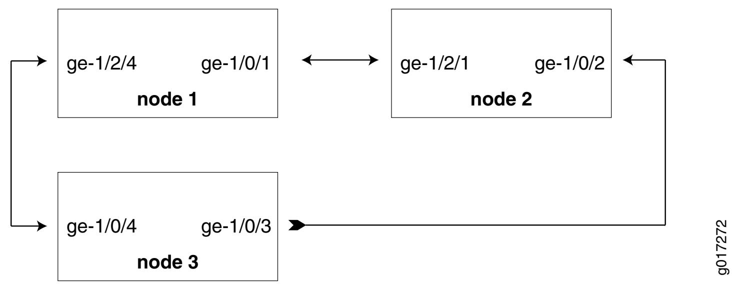

The configuration in this section is only for the RAPS channel. The bridge domain for user traffic is the same as the normal bridge domain. The only exception is if a bridge domain includes a ring port, then it must also include the other ring port of the same ring.

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 122

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;
```


Copyright © 2015, Juniper Networks, Inc.

```
        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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
                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;
              }
            }
          }
        }
      }
    }
  }
}
```

 Copyright © 2015, Juniper Networks, Inc.

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 121. 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**

Copyright © 2015, Juniper Networks, Inc.

```
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 115
- Understanding Ethernet Ring Protection Switching Functionality on page 116
- Configuring Ethernet Ring Protection Switching on page 121
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 9

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 132
- Disabling MAC Address Learning of Neighbors Through ARP or Neighbor Discovery for IPv4 and IPv6 Neighbors on page 132

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring MAC Address Validation on Static Ethernet Interfaces

To configure MAC address validation on static Ethernet interfaces, include the **mac-validate (loose | strict)** statement in the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy:

```
[edit interfaces interface-name unit logical-unit-number family family]
mac-validate (loose | strict);
```

Example of Strict MAC Validation on a Static Ethernet Interface

This example shows strict MAC address validation on a static Ethernet interface without VLAN tagging.

```
[edit interfaces]
ge-2/1/9 {
   unit 0 {
      proxy-arp;
      family inet {
         mac-validate strict;
         address 88.22.100.1/24 {
            arp 88.22.100.3 mac 00:00:58:16:64:03;
         }
      }
   }
}
```

Related Documentation

- family on page 681
- mac-validate on page 752
- *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;
```


Copyright © 2015, Juniper Networks, Inc.

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*

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 10

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 138
- 802.1Q VLAN IDs and Ethernet Interface Types on page 137
- Enabling VLAN Tagging on page 138
- Binding VLAN IDs to Logical Interfaces on page 142
- Configuring VLAN Encapsulation on page 148
- Configuring Extended VLAN Encapsulation on page 150
- Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs on page 162
- Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 151
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 151
- Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 151
- Specifying the Interface to Handle Traffic for a CCC on page 152
- Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 160
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 159
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 160
- Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 152
- Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 161
- Configuring a Logical Interface for Access Mode on page 154
- Configuring a Logical Interface for Trunk Mode on page 155
- Configuring the VLAN ID List for a Trunk Interface on page 155
- Configuring a Trunk Interface on a Bridge Network on page 156
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 5 on page 137 lists VLAN ID range by interface type.

Table 5: 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
Management and internal Ethernet interfaces	1 through 1023

Copyright © 2015, Juniper Networks, Inc.

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 for Routing Devices*.

Related Documentation

- 802.1Q VLANs Overview on page 136
- *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 Subscriber Management and Services Library*.

Related Documentation

- 802.1Q VLANs Overview on page 136
- *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. For more information, see the following sections:

Copyright © 2015, Juniper Networks, Inc.

NOTE: If you configure VLAN tagging on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces on M320, M120, and T Series routers, the 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.

Configuring Single-Tag Framing

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]
vlan-tagging;
```

Configuring Dual Tagging

To configure the routing platform 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]
stacked-vlan-tagging;
```

Configuring Mixed Tagging

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 MX Series router Gigabit and 10-Gigabit Ethernet interfaces, and for aggregated Ethernet interfaces with member links in IQ2 and IQ2-E PICs or in MX Series DPCs. Mixed tagging lets you 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 or on J Series Services Routers.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces ge-*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 ge-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
   vlan-id number;
```

```
    family family {
      address address;
    }
  }
  unit logical-unit-number {
    vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    family family {
      address address;
    }
  }
```

NOTE: When you configure the physical interface MTU for mixed tagging, 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. See "Configuring VLAN Encapsulation" on page 148.

Configuring Mixed Tagging Support for Untagged Packets

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.

Example: Configuring Mixed Tagging

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

 Copyright © 2015, Juniper Networks, Inc.

```
[edit interfaces ge-3/0/1]
flexible-vlan-tagging;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

For information about binding VLAN IDs to logical interfaces, see “Binding VLAN IDs to Logical Interfaces” on page 142. For information about configuring dual VLAN tags using the **vlan-tag** statement, see “Stacking a VLAN Tag” on page 335.

Example: Configuring Mixed Tagging to Support Untagged Packets

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 136
- *Ethernet Interfaces Feature Guide for Routing Devices*

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

Copyright © 2015, Juniper Networks, Inc.

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 138

Binding VLAN IDs to Logical Interfaces

The following sections describe how to configure logical interfaces to receive and forward VLAN-tagged frames:

Binding VLAN IDs to Logical Interfaces Overview

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 6 on page 142 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:

Table 6: 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***]**;

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***]**

Copyright © 2015, Juniper Networks, Inc.

NOTE: The **inner-list** option of the **vlan-tags** statement does not support Tag Protocol ID (TPID) values.

Binding a VLAN ID to a Logical Interface

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.

Binding a VLAN ID to a Single-Tag Logical Interface

To bind a VLAN ID to a single-tag logical interface, include the **vlan-id** statement:

vlan-id *vlan-id*;

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

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.

Binding a VLAN ID to a Dual-Tag Logical Interface

To bind a VLAN ID to a dual-tag logical interface, include the **vlan-tags** statement:

vlan-tags inner <*tpid*.>*vlan-id* outer <*tpid*.>*vlan-id*;

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

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.

Binding a Range of VLAN IDs to a Logical Interface

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.

Binding a Range of VLAN IDs to a Single-Tag Logical Interface

To bind a range of VLAN IDs to a single-tag logical interface, include the **vlan-id-range** statement:

vlan-id-range *vlan-id-vlan-id*;

Copyright © 2015, Juniper Networks, Inc.

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

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.

Binding a Range of VLAN IDs to a Dual-Tag Logical Interface

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.

```
vlan-tags inner–list [ vlan-id vlan-id–vlan-id ] outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

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.

Example: Binding Ranges VLAN IDs to Logical Interfaces

The following example configures two different ranges of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-3/0/0 {
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 500-600;
   }
}
ge-3/0/1 {
   flexible-vlan-tagging;
      unit 0 {
         encapsulation vlan-bridge;
         vlan-id-range 200-300;
      }
      unit 1 {
         encapsulation vlan-bridge;
         vlan-tags outer 1000 inner-range 100-200;
      }
}
```


Copyright © 2015, Juniper Networks, Inc.

Binding a List of VLAN IDs to a Logical Interface

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.

Binding a List of VLAN IDs to a Single-Tag Logical Interface

To bind a list of VLAN IDs to a single-tag logical interface, include the **vlan-id-list** statement. 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.

vlan-id-list [*vlan-id vlan-id–vlan-id*];

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

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.

Binding a List of VLAN IDs to a Dual-Tag Logical Interface

To bind a list of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. 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:

vlan-tags inner-list [*vlan-id vlan-id–vlan-id*] outer <*tpid*>*vlan-id*;

Copyright © 2015, Juniper Networks, Inc.

NOTE: The **inner-list** option of the **vlan-tags** statement does not support Tag Protocol ID (TPID) values.

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

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.

Example: Binding Lists of VLAN IDs to Logical Interfaces

The following example 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

Copyright © 2015, Juniper Networks, Inc.

members and apply the encapsulation type flexible-ethernet-services to that defined interface range. For more information about interface ranges, see *Configuring Interface Ranges*.

Related Documentation

- 802.1Q VLANs Overview on page 136
- *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 7 on page 147 shows the configuration statements you use to associate VLAN IDs to VLAN demux interfaces, depending on the VLAN tag framing you use:

Table 7: 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.

Copyright © 2015, Juniper Networks, Inc.

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.

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.

Configuring VLAN Encapsulation

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 virtual private LAN service (VPLS) encapsulation.

Aggregated Ethernet interfaces configured for VPLS can use Ethernet VPLS or VLAN VPLS.

To configure the encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **flexible-ethernet-services**, **vlan-ccc**, or **vlan-vpls**:

[edit interfaces *interface-name*]
encapsulation (flexible-ethernet-services | vlan-ccc | vlan-vpls);

To configure the encapsulation on an aggregated Ethernet interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **flexible-ethernet-services**, **ethernet-vpls**, or **vlan-vpls**:

[edit interfaces *interface-name*]
encapsulation (flexible-ethernet-services | ethernet-vpls | 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.

 Copyright © 2015, Juniper Networks, Inc.

In general, you configure an interface's encapsulation at the **[edit interfaces *interface-name*]** hierarchy level. However, for some encapsulation types, including 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:

```
encapsulation (vlan-ccc | vlan-tcc | vlan-vpls);
```

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*]**

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.

Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 600;
  }
}
```

Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

```
interfaces ae0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    vlan-id 100;
  }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 136
- *Configuring VPLS Interface Encapsulation*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Configuring Extended VLAN Encapsulation

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 VLAN VPLS, which allow 802.1Q tagging. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **extended-vlan-ccc** or **extended-vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | extended-vlan-vpls);
```

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.

NOTE: 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.

Example: Configuring Extended VLAN Encapsulation on a Gigabit Ethernet Interface

Configure extended VLAN CCC encapsulation on Gigabit Ethernet ingress and egress interfaces:

```
interfaces ge-0/0/0 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 {
    vlan-id 2;
    family ccc;
  }
}
interfaces ge-1/0/0 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 {
    vlan-id 2;
    family ccc;
  }
}
```

Example: Configuring Extended VLAN Encapsulation on an Aggregated Ethernet Interface

Configure extended VLAN VPLS encapsulation on an aggregated Ethernet interface:

```
interfaces ae0 {
  vlan-tagging;
  encapsulation extended-vlan-vpls;
  unit 0 {
    vlan-id 100;
  }
}
```


Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- 802.1Q VLANs Overview on page 136
- *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

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*]**

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*]**

Copyright © 2015, Juniper Networks, Inc.

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.

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
            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 136
- *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 136
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Logical Interface for Access Mode

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.

interface-mode access;

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]

 Copyright © 2015, Juniper Networks, Inc.

When an untagged or tagged packet is received on an access interface, the packet is accepted, the VLAN ID is added to the packet, and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

Example: Configuring a Logical Interface for Access Mode

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 136
- *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 136
- *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 136
- *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;
        }
    }
```

 Copyright © 2015, Juniper Networks, Inc.

```
}
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 136
- interface-mode on page 720
- *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
    }

Copyright © 2015, Juniper Networks, Inc.

```
    ...
  }
}
```

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 136
- *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 136
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 160
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 160
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 160

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.

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;
    }
```

 Copyright © 2015, Juniper Networks, Inc.

```
  }
}
```

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.

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 136
- *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 on page 163
- Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs on page 163

Copyright © 2015, Juniper Networks, Inc.

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 8 on page 163 describes the logical link-layer encapsulation types used within circuits connected using VLAN-bundled logical interfaces of the same type.

Table 8: 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.

Copyright © 2015, Juniper Networks, Inc.

Table 8: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces *(continued)*

Encapsulation Inside the Circuit	Layer 2 Circuit Joined by Configuring an Interface-to-Interface CCC Connection	
	Layer 2 VPN Routing Instance	Layer 2 Circuit
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 136
- Binding VLAN IDs to Logical Interfaces on page 142
- *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

Copyright © 2015, Juniper Networks, Inc.

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 136
- Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 152
- *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 136
- Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 161
- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 159
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 11

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 167
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Layer 2 Bridging Interfaces

You can configure an IRB logical interface at the **[edit interfaces ge-*fpc /pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
unit logical-unit-number {
}
```

You can configure Layer 3 information on the IRB logical interface by including the **irb** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
irb {
  unit logical-unit-number {
    family inet {
      address address {
      }
    }
  }
}
```

For examples of Layer 2 bridging configuration, see the *Junos OS Routing Protocols Library for Routing Devices*.

Example: Configuring Layer 2 Bridging Interfaces

The following example configures an IRB logical interface and Layer 3 information on the interface.

```
[edit interfaces]
ge-1/0/0 {
  unit 0 {
  }
}
irb {
  unit 0 {
    family inet {
    address 192.168.12.1/28;
    }
  }
}
```

Related Documentation

- family on page 681
- unit on page 892
- Layer 2 Bridging Interfaces Overview on page 167
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

 Copyright © 2015, Juniper Networks, Inc.

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 12 on page 170 shows the sample topology.

Figure 12: Configuring the MAC Address of an IRB Interface

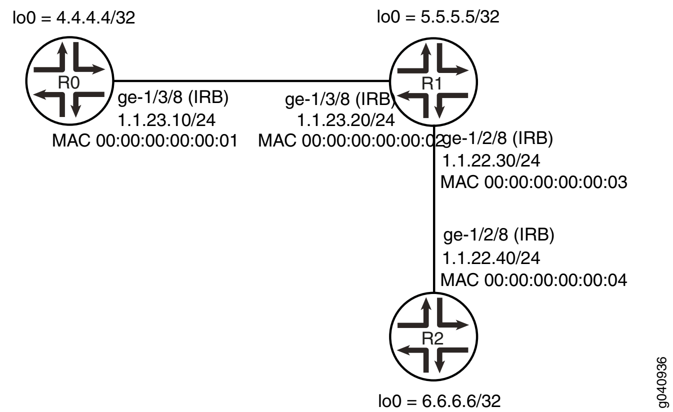

In this example you configure MAC address of IRB logical interfaces.

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
```

 Copyright © 2015, Juniper Networks, Inc.

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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

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 174

Verifying the MAC Address of the IRB Interface

Purpose Verify that the specified MAC address is assigned to the IRB interface.

 Copyright © 2015, Juniper Networks, Inc.

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 749
- *Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview*

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 12

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.

Copyright © 2015, Juniper Networks, Inc.

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 178
- Tracing LLDP Operations on page 182
- Example: Configuring LLDP on page 180

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;
  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.

 Copyright © 2015, Juniper Networks, Inc.

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.

To disable LLDP on all or a particular interface, include the **interfaces** statement at the **[edit protocols lldp]** hierarchy level:

```
interface (all | interface-name) {
  disable;
}
```

To disable LLDP on all interfaces, use the **all** option. To disable LLDP on a particular interface, include the **disable** statement with the interface name.

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*

NOTE: **The interface-name must be the physical interface (for example, ge-1/0/0) and not a logical interface (unit).**

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

Copyright © 2015, Juniper Networks, Inc.

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 by including the **interface-name** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level. When configure the **interface-name**, statement on the remote LLDP neighbor, the **show lldp neighbors** command 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.

Related Documentation

- LLDP Overview on page 177
- Tracing LLDP Operations on page 182
- Example: Configuring LLDP on page 180

- *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;
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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 177
- Configuring LLDP on page 178
- Tracing LLDP Operations on page 182

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 for Routing Devices*.

Related Documentation

- LLDP Overview on page 177
- Configuring LLDP on page 178
- Example: Configuring LLDP on page 180

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 13

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 184

- *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;
```

 Copyright © 2015, Juniper Networks, Inc.

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 680
- traceoptions on page 876
- failover-delay on page 680
- *vrrp-group*
- VRRP and VRRP for IPv6 Overview on page 183
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 14

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.

A J Series router can be configured as the CPE device for PPPoE connections. To use PPPoE, you must configure the router as a PPPoE client, encapsulate PPP packets over Ethernet, and initiate a PPPoE session.

NOTE: J4300 and J6300 routers with asymmetrical DSL (ADSL) Physical Interface Modules (PIMs) and symmetrical high-speed DSL (SHDSL) PIMs can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections. For information about configuring ADSL and SHDSL interfaces, see *ATM-over-ADSL Overview* and *ATM-over-SHDSL Overview*.

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.

On the J Series router, PPPoE establishes a point-to-point connection between the client (the Services Router) and the server, also called an access concentrator. 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:

 Copyright © 2015, Juniper Networks, Inc.

PPPoE Interfaces

The PPPoE interface to the access concentrator can be a Fast Ethernet interface on any Services Router, a Gigabit Ethernet interface on J4350 and J6350 Services Routers, an ATM-over-ADSL or ATM-over-SHDSL interface on all J Series Services Routers except the J2300, or an ATM-over-SHDSL interface on a J2300 Services Router. 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 13 on page 189 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 13: PPPoE Session on an Ethernet Loop

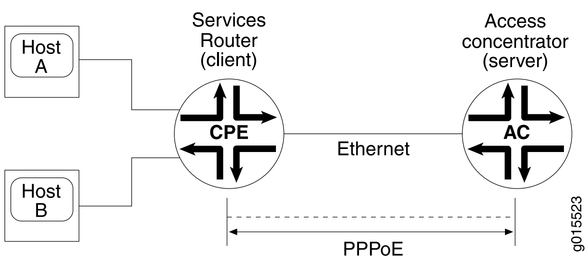

PPPoE Stages

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.

Copyright © 2015, Juniper Networks, Inc.

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.

 Copyright © 2015, Juniper Networks, Inc.

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*
- *Developing a Log Storage Strategy*
- Evaluation Order for Matching Client Information in PPPoE Service Name Tables on page 197
- Benefits of Configuring PPPoE Service Name Tables on page 198
- Configuring PPPoE on page 199
- Disabling the Sending of PPPoE Keepalive Messages on page 205
- Configuring PPPoE Service Name Tables on page 206
- Creating a Service Name Table on page 207
- Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag on page 208
- Configuring the Action Taken for the Any Service on page 209
- Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag on page 210
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 211
- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 212
- Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client on page 213
- Enabling Advertisement of Named Services in PADO Control Packets on page 214
- Assigning a Service Name Table to a PPPoE Underlying Interface on page 214
- Example: Configuring a PPPoE Service Name Table on page 216

Copyright © 2015, Juniper Networks, Inc.

- Tracing PPPoE Operations on page 221
- Troubleshooting PPPoE Service Name Tables on page 218
- Verifying a PPPoE Configuration on page 220
- *Ethernet Interfaces Feature Guide for Routing Devices*

Understanding PPPoE Service Name Tables

On an M120 router, M320 router, or MX Series router acting as a remote access concentrator (AC), also referred to as a *PPPoE server*, you can configure up to 32 PPPoE service name tables and assign the service name tables to PPPoE underlying interfaces. A *PPPoE service name table* defines the set of *services* that the router can provide to a PPPoE client. Service entries configured in a PPPoE service name table represent the *service name tags* transmitted between the client and the router in a PPPoE control packet.

This overview covers the following topics to help you understand and configure PPPoE service name tables:

Interaction Among PPPoE Clients and Routers During the Discovery Stage

In networks with mesh topologies, PPPoE clients are often connected to multiple PPPoE servers (remote ACs). During the PPPoE discovery stage, a PPPoE client identifies the Ethernet MAC address of the remote AC that can service its request, and establishes a unique PPPoE session identifier for a connection to that AC.

The following steps describe, at a high level, how the PPPoE client and the remote AC (router) use the PPPoE service name table to interact during the PPPoE discovery stage:

1. The PPPoE client broadcasts a PPPoE Active Discovery Initiation (PADI) control packet to all remote ACs in the network to request that an AC support certain services.

   The PADI packet must contain either, but not both, of the following:

   - One and only one nonzero-length service name tag that represents a specific client service

 Copyright © 2015, Juniper Networks, Inc.

- One and only one empty (zero-length) service name tag that represents an unspecified service

2. One or more remote ACs respond to the PADI packet by sending a PPPoE Active Discovery Offer (PADO) packet to the client, indicating that the AC can service the client request.

   To determine whether it can service a particular client request, the router matches the service name tag received in the PADI packet against the service name tags configured in its service name table. If a matching service name tag is found in the PPPoE service name table, the router sends the client a PADO packet that includes the name of the AC from which it was sent. If no matching service name tag is found in the PPPoE service name table, the router drops the PADI request and does not send a PADO response to the client.

3. The PPPoE client sends a unicast PPPoE Active Discovery Request (PADR) packet to the AC to which it wants to connect, based on the responses received in the PADO packets.

4. The selected AC sends a PPPoE Active Discovery Session (PADS) packet to establish the PPPoE connection with the client.

Service Entries and Actions in PPPoE Service Name Tables

A PPPoE service name table can include three types of service entries: named services, an **empty** service, and an **any** service. For each service entry, you specify the action to be taken by the underlying interface when the router receives a PADI packet containing the specified service name tag.

You can configure the following services and actions in a PPPoE service name table:

- Named service—Specifies a PPPoE client service that an AC can support. For example, you might configure named services associated with different subscribers who log in to the PPPoE server, such as **user1-service** or **user2-service**, or that correspond to different ISP service level agreements, such as **premium** and **standard**. Each PPPoE service name table can include a maximum of 512 named service entries, excluding **empty** and **any** service entries. A named service is associated with the **terminate** action by default.
- **empty** service—A service tag of zero length that represents an unspecified service. Each PPPoE service name table includes one empty service. The **empty** service is associated with the **terminate** action by default.
- **any** service—Acts as a default service for non-empty service entries that do not match the named service entries or **empty** service entry configured in the PPPoE service name table. Each PPPoE service name table includes one **any** service. The **any** service is useful when you want to match the agent circuit identifier and agent remote identifier information for a PPPoE client, but do not care about the contents of the service name tag transmitted in the control packet. The **any** service is associated with the **drop** action by default.
- Action—Specifies the action taken by the underlying PPPoE interface assigned to the PPPoE service name table on receipt of a PADI packet from the client containing a

Copyright © 2015, Juniper Networks, Inc.

particular service request. You can configure one of the following actions for the associated named service, **empty** service, **any** service, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in the PPPoE service name table on the router:

- **terminate**—(Default) Directs the router to immediately respond to the PADI packet by sending the client a PADO packet containing the name of the AC that can service the request. Named services, **empty** services, and ACI/ARI pairs are associated with the **terminate** action by default. Configuring the **terminate** action for a service enables you to more tightly control which PPPoE clients can access and receive services from a particular PPPoE server.
- **delay**—Number of seconds that the PPPoE underlying interface waits after receiving a PADI packet from the client before sending a PADO packet in response. In networks with mesh topologies, you might want to designate a primary PPPoE server and a backup PPPoE server for handling a particular service request. In such a scenario, you can configure a delay for the associated service entry on the backup PPPoE server to allow sufficient time for the primary PPPoE server to respond to the client with a PADO packet. If the primary server does not send the PADO packet within the delay period configured on the backup server, then the backup server sends the PADO packet after the delay period expires.
- **drop**—Directs the router to drop (ignore) a PADI packet containing the specified service name tag when received from a PPPoE client, which effectively denies the client's request to provide the associated service. The **any** service is associated with the **drop** action by default. To prohibit the router from responding to PADI packets that contain **empty** or **any** service name tags, you can configure the **drop** action for the empty or **any** service. You can also use the **drop** action in combination with ACI/ARI pairs to accept specific service name tags only from specific subscribers, as described in the following information about ACI/ARI pairs.

ACI/ARI Pairs in PPPoE Service Name Tables

To specify agent circuit identifier (ACI) and agent remote identifier (ARI) information for a named service, **empty** service, or **any** service in a PPPoE service name table, you can configure an ACI/ARI pair. An ACI/ARI pair contains an agent circuit ID string that identifies the DSLAM interface that initiated the service request, and an agent remote ID string that identifies the subscriber on the DSLAM interface that initiated the service request. You can think of an ACI/ARI pair as the representation of one or more PPPoE clients accessing the router by means of the PPPoE service name table.

ACI/ARI specifications support the use of wildcard characters in certain formats. You can configure a combined maximum of 8000 ACI/ARI pairs, both with and without wildcards, per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the service entries in the service name table.

You must specify the action—**terminate**, **delay**, or **drop**—taken by the underlying PPPoE interface when it receives a client request containing vendor-specific ACI/ARI information that matches the ACI/ARI information configured in the PPPoE service name table on the router. An ACI/ARI pair is associated with the **terminate** action by default.

 Copyright © 2015, Juniper Networks, Inc.

For example, assume that for the **user1-service** named service, you configure the **drop** action for the service and the **terminate** action for the associated ACI/ARI pairs. In this case, the ACI/ARI pairs identify the DSLAM interfaces and associated subscribers authorized to access the PPPoE server. Using this configuration causes the router to drop PADI packets containing the **user1-service** tag *unless* the PADI packet also contains vendor-specific ACI/ARI information that matches the subscribers identified in one or more of the ACI/ARI pairs. For PADI packets containing matching ACI/ARI information, the router sends an immediate PADO response to the client indicating that it can provide the requested service for the specified subscribers.

You can also associate a PPPoE dynamic profile, routing instance, and static PPPoE interface with an ACI/ARI pair.

Dynamic Profiles and Routing Instances in PPPoE Service Name Tables

You can associate a previously configured PPPoE dynamic profile with a named service, **empty** service, or **any** service in the PPPoE service name table, or with an ACI/ARI pair defined for these services. The router uses the attributes defined in the profile to instantiate a dynamic PPPoE subscriber interface based on the service name, ACI, and ARI information provided by the PPPoE client during PPPoE negotiation. The dynamic profile configured for a service entry or ACI/ARI pair in a PPPoE service name table overrides the dynamic profile assigned to the PPPoE underlying interface on which the dynamic PPPoE interface is created.

To specify the routing instance in which to instantiate the dynamic PPPoE interface, you can associate a previously configured routing instance with a named service, **empty** service, or **any** service in the PPPoE service name table, or with an ACI/ARI pair defined for these services. Like dynamic profiles configured for service entries or ACI/ARI pairs, the routing instance configured for the PPPoE service name table overrides the routing instance assigned to the PPPoE underlying interface.

For information about configuring the PPPoE service name table to create a dynamic PPPoE subscriber interface, see *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*.

Maximum Sessions Limit in PPPoE Service Name Tables

To limit the number of PPPoE client sessions that can use a particular service entry in the PPPoE service name table, you can configure the maximum number of active PPPoE sessions using either dynamically-created or statically-created PPPoE interfaces that the router can establish with a particular named service, **empty** service, or **any** service. (You cannot configure the maximum sessions limit for an ACI/ARI pair.) The maximum sessions limit must be in the range 1 through the platform-specific maximum PPPoE sessions supported for your routing platform. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

Copyright © 2015, Juniper Networks, Inc.

The router uses the maximum sessions value for a service entry in the PPPoE service name table in conjunction with both of the following:

- The maximum sessions (**max-sessions**) value configured for the PPPoE underlying interface
- The maximum number of PPPoE sessions supported on your routing platform

If your configuration exceeds either of these maximum session limits, the router cannot establish the PPPoE session.

Static PPPoE Interfaces in PPPoE Service Name Tables

To reserve a previously configured static PPPoE interface for use only by the PPPoE client with matching ACI/ARI information, you can specify a single static PPPoE interface for each ACI/ARI pair defined for a named service entry, **empty** service entry, or **any** service entry in a PPPoE service name table. (You cannot configure a static interface for a service entry that does not have an ACI/ARI pair defined.) The static PPPoE interface associated with an ACI/ARI pair takes precedence over the general pool of static PPPoE interfaces associated with the PPPoE underlying interface configured on the router.

When you configure a static interface in the PPPoE service name table, make sure there is a one-to-one correspondence between the PPPoE client and the static interface. For example, if two clients have identical ACI/ARI information that matches the information in the PPPoE service name table, the router reserves the static interface for exclusive use by the first client that logs in to the router. As a result, the router prevents the second client from logging in.

NOTE: **You cannot configure a static interface for an ACI/ARI pair already configured with a dynamic profile and routing instance. Conversely, you cannot configure a dynamic profile and routing instance for an ACI/ARI pair already configured with a static interface.**

PADO Advertisement of Named Services in PPPoE Service Name Tables

By default, the advertisement of named services in PADO control packets sent by the router to the PPPoE client is disabled. You can enable advertisement of named services in the PADO packet as a global option when you configure the PPPoE protocol on the router. Configuring PADO advertisement notifies PPPoE clients of the services that the router (AC) can offer.

If you enable advertisement of named services in PADO packets, make sure the number and length of all advertised service entries does not exceed the maximum transmission unit (MTU) size supported by the PPPoE underlying interface.

Related Documentation

- Evaluation Order for Matching Client Information in PPPoE Service Name Tables on page 197
- Benefits of Configuring PPPoE Service Name Tables on page 198
- Configuring PPPoE Service Name Tables on page 206

Copyright © 2015, Juniper Networks, Inc.

- Example: Configuring a PPPoE Service Name Table on page 216
- *Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles*
- PPPoE Overview on page 188

Evaluation Order for Matching Client Information in PPPoE Service Name Tables

When the router receives a service request from a PPPoE client, it evaluates the entries configured in the PPPoE service name table to find a match for the client's ACI/ARI information so it can take the appropriate action.

The order of evaluation is as follows:

1. The router evaluates the ACI/ARI information configured for the **any** service entry, and ignores the contents of the service name tag transmitted by the client.
2. If no match is found for the client information, the router evaluates the ACI/ARI information for the **empty** service entry and the named service entries. If an ACI/ARI pair is not configured for these service entries, the router evaluates the other attributes configured for the **empty** service and named services.
3. If there is still no match for the client information, the router evaluates the other attributes configured for the **any** service entry, and ignores both the ACI/ARI information for the **any** service and the contents of the service name tag transmitted by the client. If the **any** service is configured for the default action, **drop**, the router drops the PADR packet. If the **any** service is configured for a nondefault action (**terminate** or **delay**), the router evaluates the other attributes configured for the **any** service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Benefits of Configuring PPPoE Service Name Tables on page 198
- Configuring PPPoE Service Name Tables on page 206
- *Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation*
- PPPoE Overview on page 188
- *Dynamic PPPoE Subscriber Interfaces over Static Underlying Interfaces Overview*

Benefits of Configuring PPPoE Service Name Tables

This topic describes the benefits of configuring PPPoE service name tables.

Configuring PPPoE service name tables provides the following benefits:

- Enables support for multiple services requested by PPPoE clients, and configuration of an action for the underlying PPPoE interface to take (**delay**, **drop**, or **terminate**) upon receipt of a PPPoE Active Discovery Initiation (PADI) packet requesting that service.
- Provides tighter control over which PPPoE clients can log in to and receive services from a particular PPPoE server.
- Provides load balancing across a set of remote access concentrators (ACs) in a mesh topology by enabling you to configure agent circuit identifier/agent remote identifier (ACI/ARI) pairs for named, **empty**, and **any** service entries to specify the appropriate AC to receive and service a particular PPPoE client request.
- Offers a more targeted approach to configuration of PPPoE sessions based on the service name and ACI/ARI information provided by the PPPoE client during PPPoE negotiation.
- Supports creation of dynamic PPPoE subscriber interfaces in a specified routing instance based on configuration of a service entry or ACI/ARI pair in the PPPoE service name table.
- Enables you to reserve a specified static PPPoE interface for use only by the PPPoE client with matching ACI/ARI information.
- Enables you to specify the maximum number of PPPoE client sessions that can use a particular service entry in the PPPoE service name table.
- Provides redundancy across a set of remote ACs in a mesh topology by enabling you to configure a primary AC and a backup AC for handling a specific service request from a PPPoE client.

  For example, on the primary AC for handling a client service, you might configure the **terminate** action for the associated service to direct the primary AC to immediately send a PPPoE Active Discovery Offer (PADO) packet in response to a PADI packet containing that service name tag. On the backup AC for the client service, you might configure the **delay** action for the associated service to specify the number of seconds the backup AC waits after receiving a PADI packet from the client before sending a PADO packet in response. If the primary AC does not send a PADO packet to the client within the delay period configured on the backup AC, then the backup AC sends the PADO packet after the delay period expires.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- Example: Configuring a PPPoE Service Name Table on page 216
- PPPoE Overview on page 188

Copyright © 2015, Juniper Networks, Inc.

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring PPPoE

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. Optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.
9. Optionally, disable the sending of PADS messages that contain certain error tags.

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:

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 a PPPoE Interface

 Copyright © 2015, Juniper Networks, Inc.

NOTE: 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.

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]**

Copyright © 2015, Juniper Networks, Inc.

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**]**

 Copyright © 2015, Juniper Networks, Inc.

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;
```

Copyright © 2015, Juniper Networks, Inc.

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 Client Interface on a J Series Services Router

Configure a PPPoE over ATM-over-ADSL interface:

```
[edit interfaces]
at-2/0/0 {
  encapsulation ethernet-over-atm;
  atm-options {
    vpi 0;
  }
  dsl-options {
    operating-mode auto;
  }
  unit 0 {
    encapsulation ppp-over-ether-over-atm-llc;
    vci 0.120;
  }
}
pp0 {
  mtu 1492;
  unit 0 {
    ppp-options {
      chap {
        access-profile A-ppp-client;
        local-name A-at-2/0/0.0;
      }
    }
    pppoe-options {
      underlying-interface at-2/0/0;
      client;
      access-concentrator ispl.com;
      service-name "video@ispl.com";
      auto-reconnect 100;
    }
    no-keepalives;
    family inet {
      negotiate-address;
      mtu 100;
    }
    family inet6 {
      negotiate-address;
```

 Copyright © 2015, Juniper Networks, Inc.

```
            mtu 200;
        }
        family mpls {
            negotiate-address;
            mtu 300;
        }
    }
}
```

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;
            }
        }
    }
}
```

Related Documentation

- PPPoE Overview on page 188
- *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 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring PPPoE Service Name Tables

To configure PPPoE service name tables:

1. Create a PPPoE service name table.

   See “Creating a Service Name Table” on page 207.

2. (Optional) Configure the action taken for the **empty** service.

   See “Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag” on page 208.

3. (Optional) Configure the action taken for the **any** service.

   See “Configuring the Action Taken for the Any Service” on page 209.

4. Assign a named service to the service name table and optionally configure the action taken for the specified service name.

   See “Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag” on page 210.

5. (Optional) Configure the action taken for an ACI/ARI pair associated with a service.

   See “Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information” on page 211.

6. (Optional) Assign a dynamic profile and routing instance to a service name or ACI/ARI pair to instantiate a dynamic PPPoE subscriber interface.

   See *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*.

7. (Optional) Limit the number of active PPPoE sessions that the router can establish with the specified service.

   See “Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name” on page 212.

8. (Optional) Assign a static PPPoE interface to an ACI/ARI pair to reserve the interface for exclusive use by the PPPoE client with matching ACI/ARI information.

   See “Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client” on page 213.

9. (Optional) Enable advertisement of named services in the PADO control packet sent by the router to the client.

   See “Enabling Advertisement of Named Services in PADO Control Packets” on page 214.

10. Assign a service name table to a PPPoE underlying interface.

Copyright © 2015, Juniper Networks, Inc.

See “Assigning a Service Name Table to a PPPoE Underlying Interface” on page 214.

11. (Optional) Configure trace options for troubleshooting the configuration.

    See “Tracing PPPoE Operations” on page 221.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Benefits of Configuring PPPoE Service Name Tables on page 198
- Example: Configuring a PPPoE Service Name Table on page 216
- PPPoE Overview on page 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

Creating a Service Name Table

You can create up to 32 PPPoE service name tables on the router. You can optionally create named services and add them to a service name table. By default, the **empty** service and the **any** service are present in each service name table.

A named service specifies a PPPoE client service that the router, functioning as an access concentrator or PPPoE server, can support. The **empty** service is a service tag of zero length that represents an unspecified service. The **any** service acts as a default service for non-empty service entries that do not match the named or **empty** service entries configured in the PPPoE service name table. Named services and the **empty** service are associated with the **terminate** action by default, and the **any** service is associated with the **drop** action by default.

To create a PPPoE service name table:

- Specify the table name.

```
[edit protocols pppoe]
user@host# set service-name-tables table1
```

Related Documentation

- Configuring PPPoE Service Name Tables on page 206
- Understanding PPPoE Service Name Tables on page 192
- PPPoE Overview on page 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag

You can configure the action taken by the PPPoE underlying interface when it receives a PADI packet that includes a zero-length (empty) service name tag. The **empty** service is present by default in every PPPoE service name table.

To indicate that it can service the client request, the interface returns a PADO packet in response to the PADI packet. By default, the interface immediately responds to the request; this is the **terminate** action. Alternatively, you can configure the **drop** action to ignore (drop) the PADI packet, or the **delay** action to set a delay between receipt of the PADI packet and transmission of the PADO packet.

(Optional) To configure the action taken for the **empty** service in response to a PADI packet from a PPPoE client:

- Specify the action.

```
[edit protocols pppoe service-name-tables table1]
user@host# set service empty drop
```

You can also accomplish the following optional tasks when you configure the **empty** service:

- Specify the agent circuit identifier (ACI) and agent remote identifier (ARI) information to determines the action taken by the PPPoE underlying interface when it receives a PADI packet with matching ACI/ARI information.
- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Limit the number of active PPPoE sessions that the router can establish with the **empty** service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 211
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*
- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 212
- PPPoE Overview on page 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

Configuring the Action Taken for the Any Service

The **any** service acts as a default service for service name tags transmitted by the client that do not match any of the service entries configured in the PPPoE service name table on the router. By configuring an action for the **any** service, you specify the action taken by the PPPoE underlying interface when it receives a PADI control packet from a client that includes a non-empty service name tag that does not match any of the named service entries or **empty** service entry in the PPPoE service name table.

Each PPPoE service name table includes one **any** service entry associated by default with the **drop** action. The **drop** action ignores a PADI packet containing a nonmatching service name tag. Alternatively, you can configure the **terminate** action to immediately respond to the PADI packet with a PADO packet, or the **delay** action to specify a delay between receipt of the PADI packet and transmission of the PADO packet.

To configure the action taken for the **any** service in response to a PADI packet from a PPPoE client:

- Specify the action.

```
[edit protocols pppoe service-name-tables table1]
user@host# set service any terminate
```

You can also accomplish the following optional tasks when you configure the **any** service:

- Specify the agent circuit identifier (ACI) and agent remote identifier (ARI) information to determine the action taken by the PPPoE underlying interface when it receives a PADI packet with matching ACI/ARI information.
- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Limit the number of active PPPoE sessions that the router can establish with the **any** service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 211
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*
- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 212
- PPPoE Overview on page 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag

You can configure a maximum of 512 named service entries, excluding **empty** and **any** service entries, across all PPPoE service name tables on the router. A named service specifies a PPPoE client service that the router, functioning as an access concentrator or PPPoE server, can support. You can optionally configure the action taken by the PPPoE underlying interface when it receives a PADI packet that includes a matching named service (service name tag).

To indicate that it can service the client request, the interface returns a PADO packet in response to the PADI packet. By default, the interface immediately responds to the request; this is the **terminate** action. Alternatively, you can configure the **drop** action to ignore (drop) the PADI packet, or the **delay** action to set a delay between receipt of the PADI packet and transmission of the PADO packet.

(Optional) To configure a named service for a PPPoE service name table, do one of the following:

- Assign a service name to the table. The **terminate** action is applied to the service by default.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service gold-service
  ```

- Specify the action taken for a service in response to a PADI packet from a PPPoE client.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service gold-service delay 25
  ```

You can also accomplish the following optional tasks when you configure a named service:

- Specify the agent circuit identifier (ACI) and agent remote identifier (ARI) information to determines the action taken by the PPPoE underlying interface when it receives a PADI packet with matching ACI/ARI information.
- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Limit the number of active PPPoE sessions that the router can establish with the specified named service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 211

- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

 Copyright © 2015, Juniper Networks, Inc.

- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 212
- PPPoE Overview on page 188

- *Ethernet Interfaces Feature Guide for Routing Devices*

Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information

You can configure up to 8000 agent circuit identifier/agent remote identifier (ACI/ARI) pairs per PPPoE service name table, distributed in any combination among the named, **empty**, and **any** service entries in the service name table. You can optionally configure the action taken by the PPPoE underlying interface when it receives a PADI packet that includes a service name tag and the vendor-specific tag with ACI/ARI information that matches the ACI/ARI pair that you specify.

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.

To indicate that it can service the client request, the interface returns a PADO packet in response to the PADI packet. By default, the interface immediately responds to the request; this is the **terminate** action. Alternatively, you can configure the **drop** action to ignore (drop) the PADI packet, or the **delay** action to set a delay between receipt of the PADI packet and transmission of the PADO packet.

To configure an ACI/ARI pair for a named, **empty**, or **any** service, do one of the following:

- Assign an ACI/ARI pair to the service name. The **terminate** action is applied to the pair by default.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service gold-service agent-specifier aci DSLAM:3/0/1/101 ari *user*
  ```

- Specify the action taken for the ACI/ARI pair in response to a PADI packet from a PPPoE client.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service any agent-specifier aci velorum-ge-2/0/3 ari westford delay
    90
  ```

  In this example, an ACI/ARI pair and the **delay** action are configured for the **any** service. Configuring an ACI/ARI pair for the **any** service is useful when you want to match the agent circuit identifier and agent remote identifier information for a specific PPPoE

Copyright © 2015, Juniper Networks, Inc.

client, but do not care about the contents of the service name tag transmitted by the client in the PADI packet.

You can also accomplish the following optional tasks when you configure an ACI/ARI pair:

- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Reserve a specified static PPPoE interface for exclusive use by the PPPoE client with match ACI/ARI information.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*
- Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client on page 213
- PPPoE Overview on page 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name

To limit the number of PPPoE client sessions that can use a particular service entry in the PPPoE service name table, you can configure the maximum number of PPPoE sessions using static or dynamic PPPoE interfaces that the router can establish with the specified named service, **empty** service, or **any** service. You cannot configure a maximum sessions limit for an ACI/ARI pair in the service name table.

The maximum sessions limit must be in the range 1 through the platform-specific maximum PPPoE sessions supported for your routing platform. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

To limit the number of PPPoE client sessions for a particular named, **empty**, or **any** service:

- Configure the maximum sessions limit for the specified service:

```
[edit protocols pppoe service-name-tables tableEast]
user@host# set service premium-service max-sessions 100
```

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- PPPoE Overview on page 188

 Copyright © 2015, Juniper Networks, Inc.

Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client

To reserve a static PPPoE interface for exclusive use by the PPPoE client with matching agent circuit identifier/agent remote identifier (ACI/ARI) information, you can assign a previously configured static PPPoE interface to an ACI/ARI pair defined for a named service entry, **empty** service entry, or **any** service entry in a PPPoE service name table. You cannot assign a static PPPoE interface directly to a service entry that does not have an ACI/ARI pair defined.

Observe the following guidelines when you configure a static PPPoE interface for an ACI/ARI pair:

- You can specify only one static PPPoE interface per ACI/ARI pair.
- If the ACI/ARI pair represents an individual PPPoE client, make sure there is a one-to-one correspondence between the client and the static PPPoE interface.
- The static interface associated with the ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.
- You cannot configure a static interface for an ACI/ARI pair already configured with a dynamic profile and routing instance. Conversely, you cannot configure a dynamic profile and routing instance for an ACI/ARI pair already configured with a static interface.

Before you begin:

- Configure the static PPPoE interface on a M120, M320, or MX Series router.

  See "Configuring PPPoE" on page 199.

To reserve a static PPPoE interface for exclusive use by the PPPoE client with matching ACI/ARI information:

- Assign a previously configured static PPPoE interface to the ACI/ARI pair defined for a named, **empty**, or **any** service entry:

  ```
  [edit protocols pppoe service-name-tables tableEast]
  user@host# set service any agent-specifier aci velorum-ge-2/0/3 ari westford
    static-interface pp0.100
  ```

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- PPPoE Overview on page 188

Enabling Advertisement of Named Services in PADO Control Packets

You can enable advertisement of named services in PADO control packets sent by the router to the PPPoE client to indicate the services that the router can offer. By default, advertisement of named services in PADO packets is disabled. You can enable PADO advertisement as a global option on the router when you configure the PPPoE protocol.

NOTE: **Make sure the combined number and length of all named services advertised in the PADO packet does not exceed the MTU size of the PPPoE underlying interface.**

To enable advertisement of named services in PADO packets:

- Configure the PPPoE protocol to enable PADO advertisement:

  [edit protocols pppoe]
  user@host# **set pado-advertise**

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- PPPoE Overview on page 188

Assigning a Service Name Table to a PPPoE Underlying Interface

You must assign the PPPoE service name table to a PPPoE underlying interface.

Before you begin:

- Specify PPPoE as the encapsulation method on the underlying interface.

  See *Setting the Appropriate Encapsulation on the PPPoE Interface* in "Configuring PPPoE" on page 199.

To assign a service name table to a PPPoE underlying interface:

- Specify the table name:

  [edit interfaces *interface-name* unit *logical-unit-number*]
  user@host# **set pppoe-underlying-options service-name-table table1**

Related Documentation

- Configuring PPPoE Service Name Tables on page 206
- Example: Configuring a PPPoE Service Name Table on page 216
- PPPoE Overview on page 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

Disabling the Sending of PPPoE Access Concentrator Tags in PADS Packets

By default, a router that functions as an access concentrator (AC) sends the AC-Name and AC-Cookie tags, along with the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags, in the PPPoE Active Discovery Session (PADS) packet when it confirms a session with a PPPoE client. The AC-Name and AC-Cookie tags are defined as follows:

- AC-Name—String that uniquely identifies the particular AC
- AC-Cookie—Tag used by the AC to help protect against denial-of-service (DoS) attacks

If it is necessary for compatibility with your network equipment, you can prevent the router from sending the AC-Name and AC-Cookie tags in the PADS packet.

To prevent the router from transmitting the AC-Name and AC-Cookie tags in the PADS messages:

- Specify that PADS messages with AC-Name and AC-Cookie tags are not sent.

  ```
  [edit protocols pppoe]
  user@host# set no-send-pads-ac-info
  ```

  The **no-send-pads-ac-info** statement affects PADS packets sent only on PPPoE interfaces configured on the router after you configure this statement. It has no effect on PADS packets sent on previously created PPPoE interfaces.

Related Documentation

- PPPoE Overview on page 188

Discarding PADR Messages to Accommodate Abnormal CPE Behavior

This topic describes how to avoid a situation where certain CPEs respond inappropriately to normal router behavior.

During PPPoE session negotiation, the router returns PADS messages in response to PADR messages when it accepts or rejects the PPPoE session. The router adds an error tag to the PADS message when it detects a problem.

AC-System-Error is one such tag. This tag is inserted when the router imposes automatic throttling in response to excessive CPU consumption, excessive subscriber connections, or physical interfaces cycling up and down.

When the CPE receives a PADS message with this tag, the typical behavior is to retry sending PADR messages to the router or to restart session negotiation by sending PADI messages. However, some CPEs may respond inappropriately with the result that their subscribers are never connected until the CPE is rebooted.

To avoid this situation when such CPEs have access to your network, you can include the **no-send-pads-error** statement at the **[edit protocols pppoe]** hierarchy level. This statement causes the router to silently discard PADR messages in situations where the PADS would include the AC-System-Error tag. The consequence is that the CPE resends

Copyright © 2015, Juniper Networks, Inc.

PADR messages. When the conditions that result in the AC-System-Error tag are no longer present, the router once again evaluates PADR packets to determine whether to accept or reject the session.

To silently discard PADR packets:

- Specify that PADS messages with AC-System-Error tags are not sent.

```
[edit protocols pppoe]
user@host# set no-send-pads-error
```

Related Documentation

- PPPoE Overview on page 188

Example: Configuring a PPPoE Service Name Table

This example shows how you can configure a PPPoE service name table on an M120 router, M320 router, or MX Series router with service entries that correspond to different client services. By configuring the appropriate actions (**delay**, **terminate**, or **drop**) and agent circuit identifier/agent remote identifier (ACI/ARI) pairs for the service entries, you can provide load balancing and redundancy across a set of remote access concentrators (ACs) in a mesh topology, and determine how best to allocate service requests from PPPoE clients to the servers in your network.

In this example, the PPPoE service name table, Table1, contains the following service entries:

- **user1-service**—Named service representing the subscriber service for user1.
- **user2-service**—Named service representing the subscriber service for user2.
- **empty** service—Represents an unspecified service.

To configure a PPPoE service name table with service entries that correspond to different subscriber services:

1. Create the PPPoE service name table and define the services and associated actions.

```
[edit protocols pppoe]
service-name-tables Table1 {
  service empty {
    drop;
  }
  service user1-service {
    terminate;
    agent-specifier {
      aci “east*” ari “wfd*” delay 10;
      aci “west*” ari “svl*” delay 10;
    }
  }
  service user2-service {
    delay 20;
  }
}
```

 Copyright © 2015, Juniper Networks, Inc.

This example creates a PPPoE service name table named Table1 with three service entries, as follows:

- The **empty** service is configured with the **drop** action. This action prohibits the router (AC) from responding to PADI packets from the client that contain empty service name tags.
- The **user1-service** named service is configured with both the **terminate** action, and two ACI/ARI (agent-specifier) pairs:
  - The **terminate** action directs the router to immediately respond to PADI packets from the client that contain the **user1-service** tag, and is the default action for named services.
  - The 10-second delay configured for each ACI/ARI pair applies only to PADI packets from the client that contains a vendor-specific tag with matching ACI and ARI information. In this example, configuring the **delay** action indicates that the **east** or **west** server is considered the backup AC for handling these client requests, and that you expect an AC other than **east** or **west** to handle the request as the primary server. If the primary AC does not respond to the client with a PADO packet within 10 seconds, then the **east** or **west** backup AC sends the PADO packet after the 10-second delay expires.
- The **user2-service** named service is configured with a 20-second delay, indicating that you expect an AC other than the one on which this PPPoE service name table is configured to be the primary AC for handling this client request. If the primary AC does not respond to the client with a PADO packet within 20 seconds, then the backup AC (that is, the router on which you are configuring the service name table) sends the PADO packet after the 20-second delay expires.

2. Assign the PPPoE service name table to a PPPoE underlying interface configured with PPPoE encapsulation.

```
[edit interfaces]
ge-2/0/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    encapsulation ppp-over-ethernet;
    pppoe-underlying-options {
      service-name-table Table1;
    }
  }
}
```

3. (Optional) Verify the PPPoE service name table configuration.

```
user@host> show pppoe service-name-tables Table1
Service Name Table: Table1
  Service Name: <empty>
  Service Action: Drop

  Service Name: user1-service
  Service Action: Terminate
    ACI: east*
    ARI: wfd*
```

```
      ACI/ARI Action: Delay 10 seconds
    ACI: west*
    ARI: svl*
      ACI/ARI Action: Delay 10 seconds

  Service Name: user2-service
  Service Action: Delay 20 seconds
```

4. (Optional) Verify whether the PPPoE service name table has been properly assigned to the underlying PPPoE interface, and whether packet transfer between the router (AC) and PPPoE client is working correctly.

```
user@host>show pppoe underlying-interfaces ge-2/0/3.0 extensive
ge-2/0/3.0 Index 72
  State: Static, Dynamic Profile: None,
  Max Sessions: 4000, Active Sessions: 2,
  Service Name Table: Table1, Duplicate Protection: Off,
  AC Name: east
  PacketType                         Sent          Received
    PADI                                0                 2
    PADO                                2                 0
    PADR                                0                 2
    PADS                                2                 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
```

Examine the command output to ensure the following:

- The **Service Name Table** field displays the name of the correct PPPoE service name table. This field displays **none** if no service name table has been associated with the specified interface.
- The **Sent** and **Received** values for the **Service name error** field are 0 (zero). For example, a nonzero value in the **Received** field for **Service name error** indicates that there are errors in the control packets received from PPPoE clients, such as a PADI packet that does not contain a service name tag.

Related Documentation

- Understanding PPPoE Service Name Tables on page 192
- Configuring PPPoE Service Name Tables on page 206
- Troubleshooting PPPoE Service Name Tables on page 218
- PPPoE Overview on page 188

- *Ethernet Interfaces Feature Guide for Routing Devices*

Troubleshooting PPPoE Service Name Tables

Problem

Description: A misconfiguration of a PPPoE service name table can prevent PPPoE services from being properly activated. Configuration options for PPPoE service name tables are simple, which should simplify discovering where a misconfiguration exists. PPPoE clients cannot connect if the service name table contains no match for the service name tag carried in the PADI packet.

Copyright © 2015, Juniper Networks, Inc.

Symptoms: The symptom of a service name table misconfiguration is that the client connection process stops at the negotiation stage and the PADI packets are ignored. You can use the **show pppoe statistics** command to examine the PPPoE packet counts for a problem.

When the service name table is properly configured, packets sent and received increment symmetrically. The following sample output shows a PADO sent count equal to the PADI received count, and PADS sent count equal to the PADR received count. This output indicates that the PPPoE negotiation is proceeding successfully and that the service name table is not misconfigured.

```
user@host> show pppoe statistics ge-2/0/3.1

 Active PPPoE sessions: 2
    PacketType                        Sent         Received
      PADI                               0               16
      PADO                              16                0
      PADR                               0               16
      PADS                              16                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
```

When the service name table is misconfigured, the output of the **show pppoe statistics** command indicates that the number of PADI packets received on the underlying interface is increasing, but the number of PADO packets sent remains at zero. The following sample output shows a PADI count of 100 and a PADO count of 0.

```
user@host> show pppoe statistics ge-2/0/3.1

Active PPPoE sessions: 0
    PacketType                        Sent         Received
      PADI                               0              100
      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
```

When you believe a misconfiguration exists, use the **monitor traffic interface** command on the underlying interface to determine which service name is being requested by the PPPoE client. The following sample output shows that the client is requesting Service1 in the service name tag.

```
user@host> monitor traffic interface ge-2/0/3.1 print-hex print-ascii
 Listening on ge-2/0/3.1, capture size 96 bytes

 11:49:41.436682  In PPPoE PADI [Service-Name "Service1"] [Host-Uniq UTF8]
[TAG-0x120 UTF8] [Vendor-Specific UTF8]
 0x0000   ffff ffff ffff 0090 1a42 0ac1 8100 029a     .........B......
 0x0010   8863 1109 0000 00c9 0101 0008 5365 7276     .c..........Serv
 0x0020   6963 6531 0103 0004 1200 9c43 0120 0002     ice1.......C....
```

```
0x0030    044a 0105 00ab 0000 0de9 0124 783a 3132     .J.........$x:12
0x0040    3030 3963                                   009c
```

You can then use the **show pppoe service-name-tables** command to determine whether you have misspelled the name of the service or perhaps not configured the service at all.

Cause Typical misconfigurations appear in the service name table configurations.

Solution Use the appropriate statements to correct the misconfiguration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 206
- **show pppoe service-name-tables on page 1421**
- **show pppoe statistics on page 1426**
- **show pppoe underlying-interfaces on page 1428**
- PPPoE Overview on page 188
- *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 188
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 *Junos OS System Log Messages Reference*.)

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

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.

Copyright © 2015, Juniper Networks, Inc.

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**.

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.

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.

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

- PPPoE Overview on page 188
- *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.

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Tracing PPPoE Operations on page 221
- traceoptions (PPPoE) on page 880

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 221
- traceoptions (PPPoE) on page 880

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 221
- traceoptions (PPPoE) on page 880

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 221
- traceoptions (PPPoE) on page 880

Copyright © 2015, Juniper Networks, Inc.

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 221

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 221
- **traceoptions (PPPoE) on page 880**

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 15

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.

Copyright © 2015, Juniper Networks, Inc.

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 14 on page 226 and Figure 15 on page 227 show examples of situations in which you might want to configure unrestricted proxy ARP.

In Figure 14 on page 226, 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 15 on page 227, 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 14: Edge Device Case for Unrestricted Proxy ARP

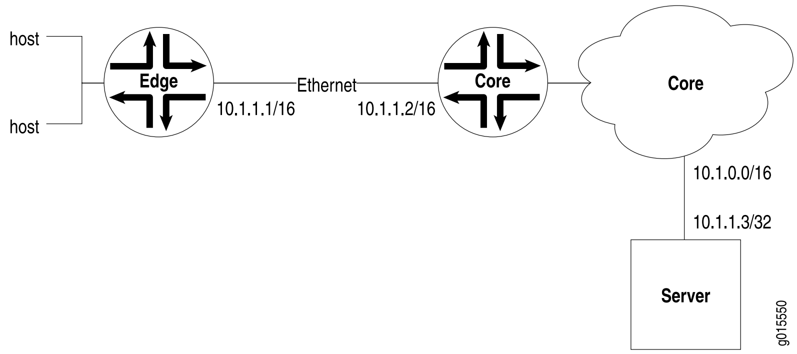

Copyright © 2015, Juniper Networks, Inc.

Figure 15: Core Device Case for Unrestricted Proxy ARP

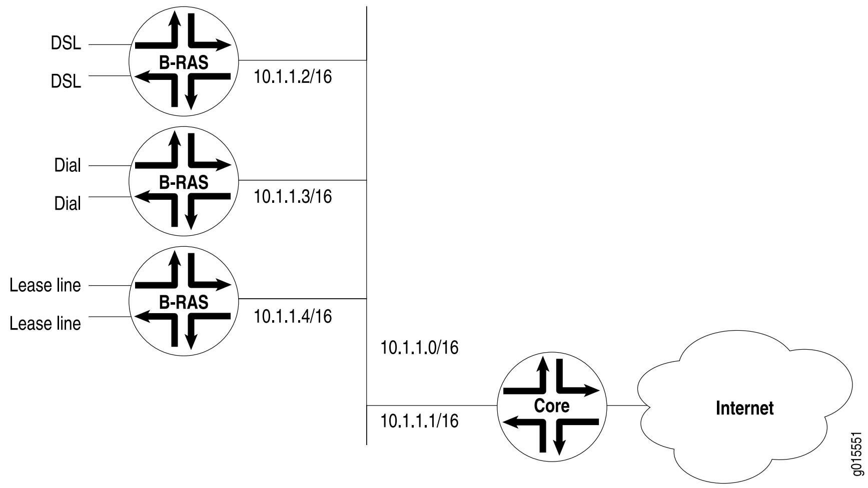

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 227
- *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 20 for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- **proxy-arp on page 828**
- Restricted and Unrestricted Proxy ARP Overview on page 225
- Configuring Gratuitous ARP on page 20

- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 16

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 on page 229

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Static ARP Table Entries

To configure static ARP table entries, include the **arp** statement:

arp *ip-address* (mac | multicast-mac) *mac-address* <publish>;

You can include this statement 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 IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, include the **multicast-mac** statement.

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**.

For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

Copyright © 2015, Juniper Networks, Inc.

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: The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

Example: Configuring Static ARP Table Entries

Configure two static ARP table entries on the router or switch's management interface:

```
[edit interfaces]
fxp0 {
  unit 0 {
    family inet {
      address 10.10.0.11/24 {
        arp 10.10.0.99 mac 0001.0002.0003;
        arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
      }
    }
  }
}
```

Related Documentation

- Management Ethernet Interface Overview on page 25
- *EX Series Switches Interfaces Overview*
- *Applying Policers*
- *Configuring an Unnumbered Interface*
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 17

Configuring TCC and Layer 2.5 Switching

TCC and Layer 2.5 Switching Overview

Translational cross-connect (TCC) is a switching concept that 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.

Related Documentation

- Configuring VLAN TCC Encapsulation on page 231
- Configuring Ethernet TCC on page 233
- *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 664
- remote on page 834
- proxy on page 827
- TCC and Layer 2.5 Switching Overview on page 231

 Copyright © 2015, Juniper Networks, Inc.

- Configuring Ethernet TCC on page 233

- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Ethernet TCC

For Layer 2.5 virtual private networks (VPNs) using an Ethernet interface as the TCC router, you can configure an Ethernet TCC.

To configure an Ethernet TCC, include the **encapsulation** statement and specify the **ethernet-tcc** option at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation ethernet-tcc;
```

For Ethernet TCC encapsulation, you must also 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.

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.

Example: Configuring an Ethernet TCC or Extended VLAN TCC

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. Ethernet TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. (See the topology in Figure 16 on page 234.)

The Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit accepts packets carrying standard TPID values.

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.

Figure 16: Topology of Layer 2.5 Translational Cross-Connect

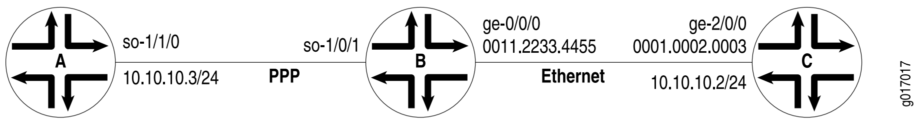

On Router B

```
interfaces
```

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. Extended VLAN TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. (See the topology in Figure 16 on page 234.)

Configuring an Extended VLAN TCC

The Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit is Ethernet with VLAN tagging enabled.

On Router B

```
interfaces
  ge-0/0/0 {
    vlan-tagging;
  encapsulation extended-vlan-tcc;
    unit 0 {
      vlan-id 1;
      family tcc {
        proxy {
          inet-address 10.10.10.3;
        }
        remote {
          inet-address 10.10.10.2;
        }
      }
    }
```

Related Documentation

- encapsulation on page 664
- remote on page 834
- proxy on page 827
- TCC and Layer 2.5 Switching Overview on page 231
- Configuring VLAN TCC Encapsulation on page 231
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

PART 2

Gigabit Ethernet Interfaces

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 18

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:

Copyright © 2015, Juniper Networks, Inc.

- **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

Copyright © 2015, Juniper Networks, Inc.

- 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 9 on page 240 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

Table 9: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs

Capability	Support
Maximum VLANs per PIC	4065
Maximum VLANs per port	1022
MAC learning per port	960
MAC accounting per port	960

 Copyright © 2015, Juniper Networks, Inc.

Table 9: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs *(continued)*

Capability	Support
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

Related Documentation

- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 245
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 246
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- *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.

Copyright © 2015, Juniper Networks, Inc.

- 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.

Copyright © 2015, Juniper Networks, Inc.

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 237

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 245.
- 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 241
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 245

 Copyright © 2015, Juniper Networks, Inc.

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 250
- mixed-rate-mode on page 766

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:

Copyright © 2015, Juniper Networks, Inc.

```
[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 237
- 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 243
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 246

- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- *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

 Copyright © 2015, Juniper Networks, Inc.

- 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 17 on page 247.

Figure 17: Control Queue Rate Limiter Scenario

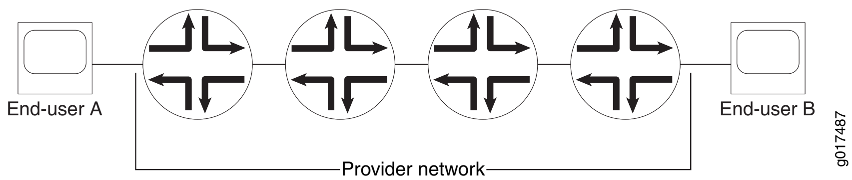

If the end users generate a mass of malicious traffic for which the port number is 179 (BGP), the router dispatches that traffic to the ingress control queue. Further, if congestion occurs in this ingress control queue due to this malicious traffic, the provider's network control packets may be affected.

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.

Copyright © 2015, Juniper Networks, Inc.

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;
  }
```


Copyright © 2015, Juniper Networks, Inc.

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 237
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 245
- **no-pre-classifier on page 777**
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC

Table 10 on page 250 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.

Copyright © 2015, Juniper Networks, Inc.

Table 10: 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 237
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 245
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 246
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Mixed-Rate Mode Operation

To configure mixed-rate mode operation in a the 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
```


Copyright © 2015, Juniper Networks, Inc.

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 245
- **mixed-rate-mode on page 766**

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 19

Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength

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 253
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **wavelength on page 914**

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;
```

Copyright © 2015, Juniper Networks, Inc.

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 11 on page 254 shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

Table 11: 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
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

Copyright © 2015, Juniper Networks, Inc.

Table 11: Wavelength-to-Frequency Conversion Matrix *(continued)*

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
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		
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 253
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **wavelength on page 914**

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 20

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*.

Copyright © 2015, Juniper Networks, Inc.

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.

Related Documentation

- *Configuring SONET/SDH Framing Mode*
- Configuring 10-Gigabit Ethernet Framing on page 259
- Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces on page 258
- *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 698**
- Configuring 10-Gigabit Ethernet Framing on page 259
- 10-Gigabit Ethernet Framing Overview on page 257
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 698**
- 10-Gigabit Ethernet Framing Overview on page 257
- Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces on page 258

- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 21

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

10-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 in M120, M320. and T Series routers.

Related Documentation

- Configuring 10-Gigabit Ethernet Notification of Link Down Alarm on page 261
- *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 262.

Related Documentation

- Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning on page 262
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

To configure this option, include the **asynchronous-notification** statement at the **[edit interfaces ge-** ***fpc/pic/port*** **gigether-options]** hierarchy level:

Copyright © 2015, Juniper Networks, Inc.

```
[edit interfaces]
ge-fpc/pic/port {
   gigether-options {
      asynchronous-notification;
   }
}
```

Related Documentation

- 10-Gigabit Ethernet Notification of Link Down Alarm Overview on page 261
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 618
- warning on page 913
- 10-Gigabit Ethernet Notification of Link Down for Optics Options Overview on page 261
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 22

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

Copyright © 2015, Juniper Networks, Inc.

    - 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 265. 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 265
- *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*

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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 a Logical Interface for Access Mode” on page 154 and “Configuring a Logical Interface for Trunk Mode” on page 155.

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 42.
   - 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 42 and “10-port 10-Gigabit Ethernet LAN/WAN PIC Overview” on page 237.

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 311, “Configuring MAC Address Filtering” on page 316, and “Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview” on page 325.

Related Documentation

- 40-Gigabit Ethernet PIC Overview on page 263
- Configuring Gigabit Ethernet Policers on page 311
- Configuring MAC Address Filtering on page 316
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 325

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 23

Configuring 100-Gigabit Ethernet PICs/MICs

100-Gigabit Ethernet Interfaces Overview

100-Gigabit Ethernet

Juniper Networks supports a variety of !00-Gigabit Ethernet interfaces. 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.

Copyright © 2015, Juniper Networks, Inc.

100-Gigabit Ethernet transmits Ethernet frames at a rates of 100 gigabits per second, and was first defined by the IEEE 802.3ba-2010 standard. It is used for very high speed transmission of voice and data signals across the numerous world-wide fiber-optic networks.

MX Series 100-Gigabit Ethernet Interfaces

Table 12 on page 268 lists the 100-Gigabit Ethernet interfaces supported by MX Series routers.

Table 12: 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 269
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 269
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*

PTX Series 100-Gigabit Ethernet Interfaces

Table 13 on page 268 lists the 100-Gigabit Ethernet interfaces supported by PTX Series routers.

Table 13: 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)*

 Copyright © 2015, Juniper Networks, Inc.

T Series 100-Gigabit Ethernet Interfaces

Table 14 on page 269 lists the 100-Gigabit Ethernet interfaces supported by T Series routers.

Table 14: 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 271 “Configuring 100-Gigabit Ethernet Type 4 PIC With CFP” on page 274
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 280

Related Documentation

- *MICs Supported by MX Series Routers*
- *MPCs Supported by MX240, MX480, MX960, MX2010, and MX2020 Routers*
- *PTX Series PICs Supported*
- *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)

Copyright © 2015, Juniper Networks, Inc.

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 MX240, MX480, MX960, MX2010, and MX2020 MPC3E*

 Copyright © 2015, Juniper Networks, Inc.

- *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.

 Copyright © 2015, Juniper Networks, Inc.

- 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 274
- *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.

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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.

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

- 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 311, "Configuring MAC Address Filtering" on page 316, and "Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview" on page 325.

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.

Related Documentation

- 100-Gigabit Ethernet Type 4 PIC with CFP Overview on page 271
- Configuring Gigabit Ethernet Policers on page 311
- Configuring MAC Address Filtering on page 316
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 325

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.

 Copyright © 2015, Juniper Networks, Inc.

- 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 693
- sa-multicast (100-Gigabit Ethernet) on page 842
- vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 906
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 907

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.

 Copyright © 2015, Juniper Networks, Inc.

  - **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 311, “Configuring MAC Address Filtering” on page 316, and “Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview” on page 325.
- 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.

Copyright © 2015, Juniper Networks, Inc.

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 15 on page 282 lists the capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP.

Table 15: 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 274
- Configuring Gigabit Ethernet Policers on page 311
- Configuring MAC Address Filtering on page 316
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 325

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

Copyright © 2015, Juniper Networks, Inc.

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 693**
- **sa-multicast on page 842**
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 285

- *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.

Copyright © 2015, Juniper Networks, Inc.

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 287
- **sa-multicast on page 843**

Copyright © 2015, Juniper Networks, Inc.

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 282

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 285
- Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface on page 286

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;
         }
       }
     }
   ```

Copyright © 2015, Juniper Networks, Inc.

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.

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]

Copyright © 2015, Juniper Networks, Inc.

```
user@host# show interfaces
  et-5/0/0:0 {
    gigether-options {
      802.3ad ae0;
    }
  }
  et-5/0/0:1 {
    gigether-options {
      802.3ad ae0;
    }
  }
```

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 282
- *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 on page 287
- Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface on page 288

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
```

 Copyright © 2015, Juniper Networks, Inc.

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;
    }
  }
```

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP on page 283
- **sa-multicast on page 843**
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 282

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 11 on page 254 shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

Table 16: 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
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

Copyright © 2015, Juniper Networks, Inc.

Table 16: Wavelength-to-Frequency Conversion Matrix *(continued)*

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
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		
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 253
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **wavelength on page 914**

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 24

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;
  }
  (prbs | no-prbs);
  preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
  }
  rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    otu4;
    (pass-through | no-pass-through);
  }
  signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
```

Copyright © 2015, Juniper Networks, Inc.

```
      interval value;
    }
    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 294
- *Ethernet Interfaces Feature Guide for Routing Devices*

10-Gigabit Ethernet OTN Options Configuration Overview

M120, 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.

Related Documentation

- otn-options on page 788
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 295 for more information.

Copyright © 2015, Juniper Networks, Inc.

For more information about the 100-Gigabit DWDM OTN PIC, see *100-Gigabit DWDM OTN PIC* in the *PTX Series Interface Module Reference*.

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, and 100-Gigabit Ethernet)**, and **show interfaces interval** 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 Ethernet OTN Optics on page 299
- *Ethernet Interfaces Feature Guide for Routing Devices*
- **show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet) on page 1171**
- **optics-options on page 787**
- **otn-options on page 788**

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

Copyright © 2015, Juniper Networks, Inc.

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 18 on page 296. 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 18: Pre-FEC BER Monitoring

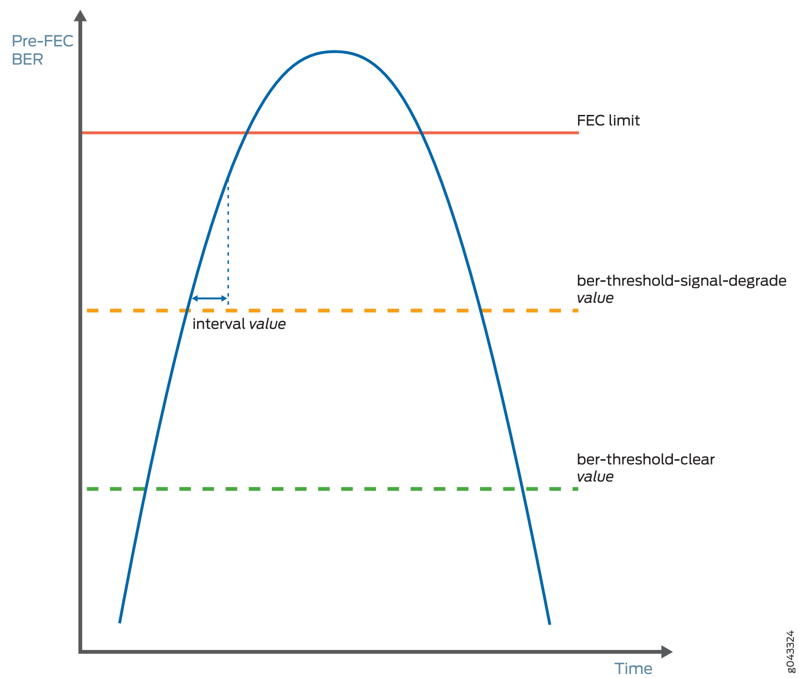

With pre-FEC BER monitoring enabled, when the configured pre-FEC BER signal degrade threshold is reached, the PIC stops forwarding packets to the remote interface and raises an interface alarm. Ingress packets continue to be processed. If pre-FEC BER monitoring is used with MPLS fast reroute or another 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

Copyright © 2015, Juniper Networks, Inc.

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 17 on page 298 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 17: 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
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 18 on page 298 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 18: 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

 Copyright © 2015, Juniper Networks, Inc.

Table 18: Example—Signal Degrade and Clear Thresholds After Configuration *(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	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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Configuring 100-Gigabit Ethernet OTN Optics

PTX Series routers support Optical Transport Network (OTN) interfaces, including the 100-Gigabit Ethernet DWDM OTN PIC.

To configure an OTN interface:

1. Configure the interface wavelength. See “Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength” on page 253.
2. Enable the laser:

   [edit interfaces et-0/0/0 otn-options]
   user@host# **set laser-enable**

3. (Optional) Set the trace identifiers:

   [edit interfaces et-0/0/0 otn-options]
   user@host# **set tti *tti-identifier tti-identifier-name***

4. (Optional) Set the triggers:

   [edit interfaces et-0/0/0 otn-options]
   user@host# **set trigger *trigger-identifier***

5. (Optional) Enable VLAN tagging. See *Enabling VLAN Tagging*.
6. (Optional) Set the media MTU. See *Configuring the Media MTU*.
7. (Optional) Set the unit VLAN ID, family inet, and IP address:

   [edit interfaces et-0/0/0]
   user@host# **set vlan-id *number***
   user@host# **set family inet**
   user@host# **set address *address***

8. (Optional) Configure pre-forward error correction (pre-FEC)-based fast reroute (FRR) to enable monitoring of the pre-FEC status of the OTN link:

   a. Set the BER signal-degrade threshold:

      [edit interfaces et-0/0/0 otn-options signal-degrade]
      user@host# **set ber-threshold-signal-degrade *value***

   b. Set the BER threshold to clear signal-degrade alarms:

      [edit interfaces et-0/0/0 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.

      [edit interfaces et-0/0/0 otn-options signal-degrade]
      user@host# **set interval *value***

      NOTE: Configuring a high BER threshold for signal-degrade and a long interval might cause the internal counter register to be saturated. An invalid configuration will be discarded by the PIC driver, and the default values will be used instead. A system log message will be logged for this error.

   d. Enable signal degrade monitoring:

      [edit interfaces et-0/0/0 otn-options preemptive-fast-reroute]
      user@host# **set signal-degrade-monitor-enable**

9. (Optional) Enable backward FRR to insert the local pre-FEC status into transmitted OTN frames and monitor received OTN frames for the pre-FEC status:

   [edit interfaces et-0/0/0 otn-options preemptive-fast-reroute]
   user@host# **set backward-frr-enable**

Related Documentation

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- optics-options on page 787
- otn-options on page 788
- signal-degrade on page 850

Copyright © 2015, Juniper Networks, Inc.

- preemptive-fast-reroute on page 825

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 25

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 19 on page 304 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 19: 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
VPLS	Yes	Yes

Copyright © 2015, Juniper Networks, Inc.

For more information about configuring VPLS, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS, Release 14.1*.

You can also configure CoS on logical IQ interfaces. For more information, see the *Class of Service Feature Guide for Routing Devices*.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 311
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 318
- Configuring MAC Address Accounting on page 305
- *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), and for Gigabit Ethernet DPCs on MX Series routers, you can configure whether source and destination MAC addresses are dynamically learned. To configure MAC address accounting, 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 prohibit the interface from dynamically learning source and destination MAC addresses, include the **no-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]
no-mac-learn-enable;
```

MAC address learning is based on source addresses. You can start accounting for traffic after it has been sent from the MAC address. Once the MAC address is learned, the frames and bytes transmitted to or received from the MAC address can be tracked.

NOTE: DPCs and MPCs support MAC address accounting. DPCs support both source and destination MAC address accounting. MPCs support only source MAC address accounting.

Related Documentation

- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 303
- Configuring Gigabit Ethernet Policers on page 311
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 318
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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), Modular Port Concentrators (MPCs), and Type-3, Type-4 and Type-5 Flexible Port Concentrators (FPCs). Also, this feature is supported on 10-Gigabit Ethernet interfaces on MX Series routers with MPC4E. 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 20 on page 306 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 20: 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)

 Copyright © 2015, Juniper Networks, Inc.

Table 20: Adjustment Bytes for Logical Interfaces over Ethernet Interfaces *(continued)*

Encapsulation Type on Logical Interfaces	Number of Adjustment Bytes	Description
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.

Copyright © 2015, Juniper Networks, Inc.

- 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.

Related Documentation

- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 308
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 309
- **account-layer2-overhead on page 604**
- *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 on page 308

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**

Copyright © 2015, Juniper Networks, Inc.

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**

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)
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
     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 306
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 308
- **show interfaces (Gigabit Ethernet) on page 1214**
- *show interfaces statistics*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gigabit Ethernet Policers

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*.

This section contains the following topics:

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

 Copyright © 2015, Juniper Networks, Inc.

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 316.

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);
    }
```

Copyright © 2015, Juniper Networks, Inc.

```
  }
}
```

You can define a forwarding class, or you can use a predefined forwarding class. Table 21 on page 314 shows the predefined forwarding classes and their associated queue assignments.

Table 21: 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*. For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 316.

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. For information about defining these policers, see "Configuring Gigabit Ethernet Policers" on page 311.

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*]

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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 314.

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 "Enabling Ethernet MAC Address Filtering" on page 15.

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;
```

For more information about the **accept-source-mac** statement, see "Applying a Policer" on page 314.

Example: Configuring Gigabit Ethernet Policers

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.

```
[edit interfaces]
ge-6/0/0 {
    gigether-options {
        ether-switch-profile {
            ether-policer-profile {
                input-priority-map {
                    ieee-802.1p {
                        premium [ 2 3 ];
```

 Copyright © 2015, Juniper Networks, Inc.

```
          }
        }
        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 303
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 318
- Configuring MAC Address Accounting on page 305
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

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 (srTCM)—A single-rate tricolor marking policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) 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). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets arriving are below the CBS (green), exceed the CBS (yellow) but not the EBS, or exceed the EBS (red). Single-rate TCM is most useful when a service is structured according to packet length and not peak arrival rate.
- Two-rate Tricolor Marking (trTCM)—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding (AF) 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 marked as belonging to one of three categories (green, yellow, or red) based on whether the packets arriving are below the CIR (green), exceed the CIR (yellow) but not the PIR, or exceed the PIR (red). Two-rate TCM is most useful when a service is structured according to arrival rates and not necessarily packet length.

Unlike policing (described in “Configuring Gigabit Ethernet Policers” on page 311), configuring two-color policers and tricolor marking policers requires that you configure a firewall filter.

 Copyright © 2015, Juniper Networks, Inc.

This section contains the following topics:

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 for Routing Devices* and the *Class of Service Feature Guide for Routing Devices*.

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;
}
```

Copyright © 2015, Juniper Networks, Inc.

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;
```

 Copyright © 2015, Juniper Networks, Inc.

```
  }
  then loss-priority-high;
}
[edit interfaces ge-1/1/0]
unit 2 {
  layer2-policer {
    input-policer two-color-policer;
    output-policer two-color-policer;
  }
}
```

Related Documentation

- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs on page 303
- Configuring Gigabit Ethernet Policers on page 311
- Configuring MAC Address Accounting on page 305
- *Configuring a Policer Overhead*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 26

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:

- For Gigabit Ethernet interfaces installed in J4350 and J6350 Services Routers, when you manually configure either the link mode or speed settings, the system ignores the configuration and generates a system log message. When autonegotiation is enabled and you specify the link mode and speed, the link autonegotiates with the manually configured settings. When autonegotiation is disabled and you configure both the link mode and speed, the link operates with the manually configured settings. If you disable autonegotiation and do not manually configure the link mode and speed, the link operates at 1000 Mbps full duplex.
- When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
- On ACX Series Universal Access 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*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 27

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 136)

Related Documentation

- 802.1Q VLANs Overview on page 136
- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 326
- *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.

 Copyright © 2015, Juniper Networks, Inc.

- **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 22 on page 327 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 22: 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

Copyright © 2015, Juniper Networks, Inc.

Table 22: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames *(continued)*

Rewrite Operation	Untagged	Single-Tagged	Dual-Tagged	Number of Tags
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 23 on page 329 shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

 Copyright © 2015, Juniper Networks, Inc.

Table 23: 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 331.

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 325
- Understanding swap-by-poppush on page 353
- **swap-by-poppush on page 861**
- *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 325
- **aggregated-ether-options on page 617**
- **ethernet-switch-profile on page 676**
- **gigether-options on page 699**
- **tag-protocol-id on page 869**
- *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 329
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 141

Copyright © 2015, Juniper Networks, Inc.

Configuring Stacked VLAN Tagging

To configure stacked VLAN tagging for all logical interfaces on a physical interface, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
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 335.

Related Documentation

- **stacked-vlan-tagging on page 857**
- Stacking a VLAN Tag on page 335
- *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 344.

Related Documentation

- **unit on page 892**
- Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 344
- *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 TPID values and inner or outer VLAN ID values. These values can be applied to either the input VLAN map or the output VLAN map.

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:

```
inner-tag-protocol-id tpid;
```

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]**

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:

```
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;
}
```

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:

```
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;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**

 Copyright © 2015, Juniper Networks, Inc.

- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-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 136.

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 329.

Table 24 on page 333 and Table 25 on page 334 specify when these statements are required. Table 24 on page 333 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 24: 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 25 on page 334 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.

Copyright © 2015, Juniper Networks, Inc.

Table 25: 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

The following examples use Table 24 on page 333 and Table 25 on page 334 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

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;
}
```

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- inner-tag-protocol-id on page 709
- input-vlan-map on page 715
- output-vlan-map on page 795
- pop-swap on page 810
- unit on page 892
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 23 on page 329 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 136.

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 331.

Related Documentation

- tag-protocol-id on page 870
- unit on page 892

- Table 23 on page 329
- 802.1Q VLANs Overview on page 136
- Configuring Inner and Outer TPIDs and VLAN IDs on page 331
- *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 24 on page 333 and Table 25 on page 334 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 715**
- **output-vlan-map on page 795**
- **pop on page 808**
- **unit on page 892**
- See Table 24 on page 333 and Table 25 on page 334 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;
```

 Copyright © 2015, Juniper Networks, Inc.

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 715
- output-vlan-map on page 795
- pop on page 808
- unit on page 892
- *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 24 on page 333 and Table 25 on page 334 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 715
- output-vlan-map on page 795
- pop-pop on page 809
- unit on page 892

- See Table 24 on page 333 and Table 25 on page 334 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 24 on page 333 and Table 25 on page 334 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 715
- output-vlan-map on page 795
- pop-swap on page 810
- unit on page 892
- See Table 24 on page 333 and Table 25 on page 334 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:

 Copyright © 2015, Juniper Networks, Inc.

```
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 136.

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 344.

Related Documentation

- **input-vlan-map on page 715**
- **output-vlan-map on page 795**
- **swap on page 861**
- **vlan-id on page 901**
- **tag-protocol-id on page 870**
- **unit on page 892**
- For more information about binding a VLAN ID to the logical interface, see 802.1Q VLANs Overview on page 136.
- For more information about the swap operation, see Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 344.

- *Ethernet Interfaces Feature Guide for Routing Devices*

Rewriting a VLAN Tag on Untagged Frames

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.

 Copyright © 2015, Juniper Networks, Inc.

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 26 on page 341 through Table 28 on page 341 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 26: 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 27: 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 28: 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
push-push	No	No	Yes

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;
  }
}
```

 Copyright © 2015, Juniper Networks, Inc.

You can use the **show interface *interface-name*** command to display the status of a modified VLAN map for the specified interface.

Related Documentation

- input-vlan-map on page 715
- output-vlan-map on page 795
- pop on page 808
- pop-pop on page 809
- push on page 823
- push-push on page 824
- unit on page 892
- *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 24 on page 333 and Table 25 on page 334 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 715
- output-vlan-map on page 795
- swap-push on page 862
- unit on page 892
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 24 on page 333 and Table 25 on page 334 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 715
- output-vlan-map on page 795
- swap-swap on page 863
- unit on page 892
- *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:

Copyright © 2015, Juniper Networks, Inc.

- Swap a TPID and VLAN ID Pair for Both VLAN Tags on Ingress and on Egress on page 349
- Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress on page 349
- Pop a TPID and VLAN ID Pair on Ingress; Push a VLAN ID and TPID Pair on Egress on page 349
- Pop an Outer VLAN Tag to Connect an Untagged VPLS Interface to Tagged VPLS Interfaces on page 350

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;
    }
```

Copyright © 2015, Juniper Networks, Inc.

```
    }
}
```

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
        ...
        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;
```

 Copyright © 2015, Juniper Networks, Inc.

```
        }
        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 Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress

```
[edit interfaces]
ge-1/1/0 {
    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;
```

Copyright © 2015, Juniper Networks, Inc.

```
            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 715
- output-vlan-map on page 795
- inner-tag-protocol-id on page 709
- inner-vlan-id on page 710
- pop on page 808
- pop-pop on page 809
- pop-swap on page 810
- push on page 823
- push-push on page 824
- swap on page 861
- swap-push on page 862
- swap-swap on page 863
- unit on page 892
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 29 on page 351 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 29: 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

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 19 on page 351 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 19: swap-push (transparent tag)

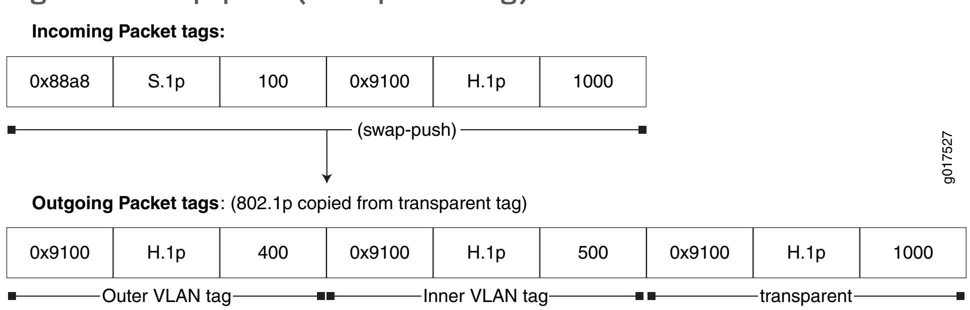

Copyright © 2015, Juniper Networks, Inc.

Figure 20 on page 352 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 20: swap-push (no transparent tag)

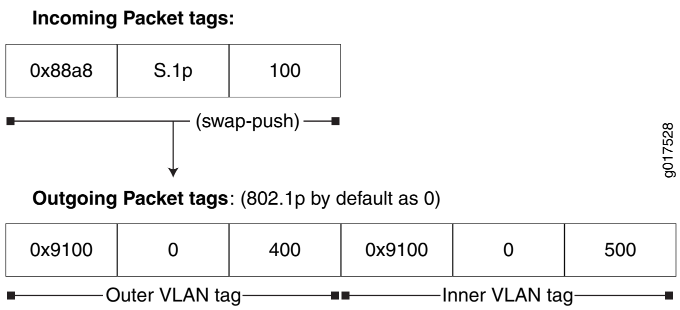

Figure 21 on page 352 shows an incoming packet with a transparent tag. A push operation pushes another VLAN tag. The IEEE 802.1p value is inherited from the transparent tag.

Figure 21: push (transparent tag)

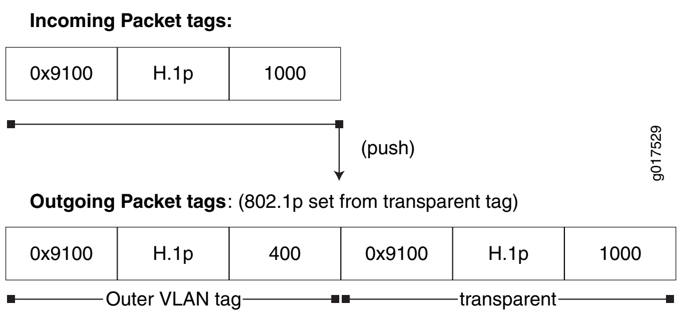

Figure 22 on page 352 shows an incoming packet with a transparent tag. A push-push operation pushes the outer and inner VLAN tags, respectively. The IEEE 802.1p values are inherited from the transparent tag.

Figure 22: push-push (transparent tag)

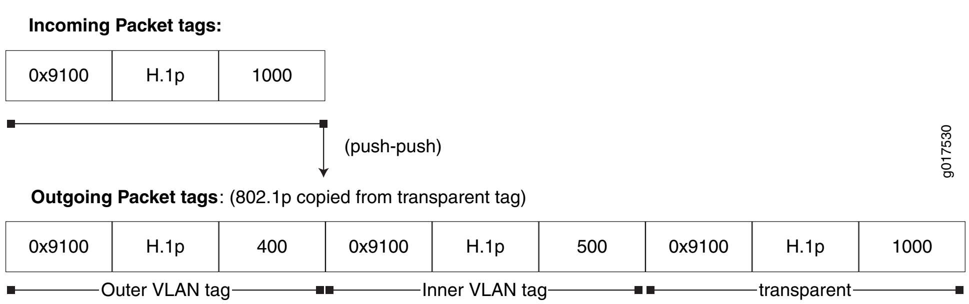

Related Documentation

- *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 353
- Understanding swap-by-poppush on page 353

Copyright © 2015, Juniper Networks, Inc.

- swap-by-poppush on page 861
- *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 861
- *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 353
- Understanding Transparent Tag Operations and IEEE 802.1p Inheritance on page 351

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;
    }
  }
}
```

 Copyright © 2015, Juniper Networks, Inc.

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 861
- *Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag*
- Understanding swap-by-poppush on page 353
- Understanding Transparent Tag Operations and IEEE 802.1p Inheritance on page 351

Copyright © 2015, Juniper Networks, Inc.

PART 3

Operation, Administration, and Management (OAM) for Ethernet Interfaces

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 28

Configuring IEEE 802.1ag OAM Connectivity-Fault Management

Copyright © 2015, Juniper Networks, Inc.

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.

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 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:

- 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.
- 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.

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.CFM does not support Multichassis Link Aggregation (MC-LAG). Do not configure the **mc-ae** statement when you configure CFM.

 Copyright © 2015, Juniper Networks, Inc.

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.

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.

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 361
- Best Practices for Configuring 802.1ag Ethernet OAM for VPLS on page 362

Connectivity Fault Management Key Elements

Figure 23 on page 361 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 23: Relationship Among MEPs, MIPs, and Maintenance Domain Levels

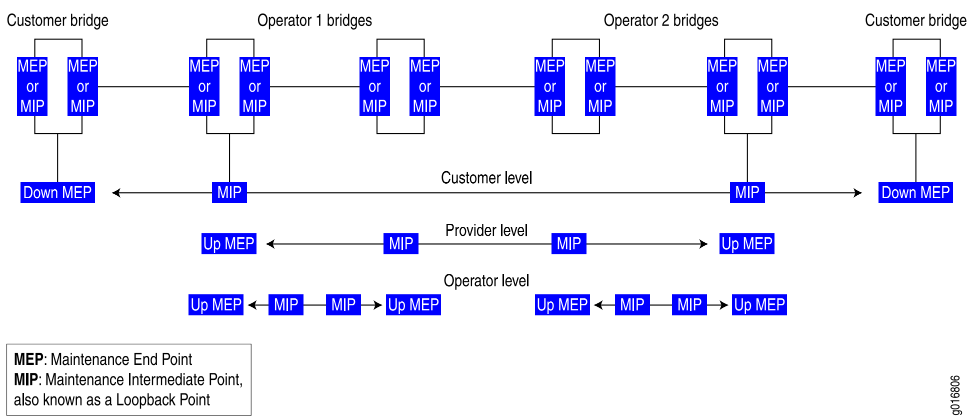

NOTE: Maintenance intermediate points (MIP) are not supported on the ACX Series routers.

Copyright © 2015, Juniper Networks, Inc.

A maintenance association is a set of MEPs configured with the same maintenance association identifier and maintenance domain level. Figure 24 on page 362 shows the hierarchical relationships between the Ethernet bridge, maintenance domains, maintenance associations, and MEPs.

Figure 24: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs

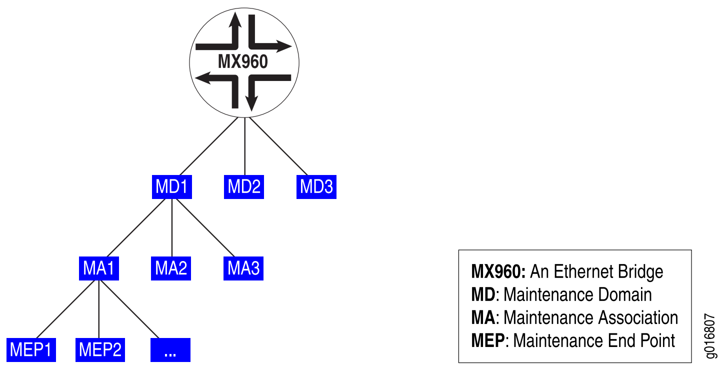

Best Practices for Configuring 802.1ag Ethernet OAM for VPLS

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 360 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:

Copyright © 2015, Juniper Networks, Inc.

- 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.

Related Documentation

- **connectivity-fault-management on page 640**
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Creating the Maintenance Domain

To enable CFM on an Ethernet interface, maintenance domains, maintenance associations, and MEPs must be created and configured.

To create a maintenance domain, include the **maintenance-domain *domain-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Copyright © 2015, Juniper Networks, Inc.

Give the maintenance domain a name. Names can be in one of several formats:

- Configuring the Maintenance Domain Name Format on page 364
- Configuring the Maintenance Domain Level on page 364

Configuring the Maintenance Domain Name Format

You can specify the maintenance domain name format as one of the following:

- A plain ASCII character string.
- A domain name service (DNS) format, a MAC address plus a two-octet identifier in the range from 0 through 65,535, or none.
- A MAC address plus a two-octet identifier in the range from 0 through 65,535.
- Or none.

If none is specified, the maintenance domain name is not used.

The default name format is an ASCII character string.

To configure the maintenance domain name format, include the **name-format (character-string | none | dns | mac+2octet)** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

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

- 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

Configuring the Maintenance Domain Level

The maintenance domain level is a mandatory parameter that indicates the nesting relationship between various maintenance domains. The level is embedded in each of the CFM frames. CFM messages within a given level are processed by MEPs at that 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.

To configure the maintenance domain level, include the **level *number*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

Related Documentation

- connectivity-fault-management on page 640
- maintenance-domain on page 754

Copyright © 2015, Juniper Networks, Inc.

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Maintenance Intermediate Points

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.

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 sections describe MIP configuration:

- Configuring MIP for Bridge Domains of a Virtual Switch on page 366
- Configuring the Maintenance Domain Bridge Domain on page 366
- Configuring the Maintenance Domain Instance on page 366
- Configuring the Maintenance Domain MIP Half Function on page 366

Configuring MIP for Bridge Domains of a Virtual Switch

The default maintenance domain configuration allows MIP configuration for bridge domains for a default virtual switch or a user-defined virtual switch. You can use the **virtual-switch** and **bridge-domain** statements to specify which MIPs to enable for a user-defined virtual switch.

A bridge domain must be specified by name only if it is configured by including the **vlan-id *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.

To configure a bridge domain under a user-defined virtual switch, include the **virtual-switch *name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* default-*x*]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  domain-name default-x]
virtual-switch name {
  bridge-domain {
    name-1;
    name-2 {
      vlan-id [vlan-ids ];
    }
  }
}
```

Configuring the Maintenance Domain Bridge Domain

The VLAN corresponds to the bridge domain.

To configure the bridge domain for the default virtual switch, include the bridge-domain statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name*]** hierarchy level.

Configuring the Maintenance Domain Instance

To configure the maintenance domain instance for a VPLS routing instance, include the instance statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level.

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

 Copyright © 2015, Juniper Networks, Inc.

MIP half function, include the mip-half-function statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level.

Related Documentation

- bridge-domain on page 627
- connectivity-fault-management on page 640
- instance on page 715
- mip-half-function on page 765
- virtual-switch on page 899
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *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 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Continuity Check Protocol

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

The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32.

- Configuring the Continuity Check on page 368
- Configuring the Continuity Check Hold Interval on page 369
- Configuring the Continuity Check Interval on page 369
- Configuring the Continuity Check Loss Threshold on page 370
- Continuity Measurement on page 370

Configuring the Continuity Check

You can configure the following continuity check protocol parameters:

- **hold-interval *minutes***
- **interval *time***
- **loss-threshold *number***

 Copyright © 2015, Juniper Networks, Inc.

To enable the continuity check protocol, include the **continuity-check** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name*****]** hierarchy level.

Configuring the Continuity Check Hold Interval

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

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 auto discovered 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.

To configure the hold interval, include the **hold-interval** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name*** **continuity-check]** hierarchy level.

NOTE: Hold timer based flushing is applicable only for auto discovered remote MEPs and not for statically configured remote MEPs.

Configuring the Continuity Check Interval

You can specify the continuity check message (CCM) interval. The interval is the time between the transmission of CCMs. You can specify 10 minutes (**10m**), 1 minute (**1m**), 10 seconds (**10s**), 1 second (**1s**), 100 milliseconds (**100ms**), or 10 milliseconds (**10ms**). The default value is 1 minute.

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 (PFE) by default. You can only disable PPM on the PFE. To disable PPM on the PFE, 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).

To configure the interval, include the **interval** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name*** **continuity-check]** hierarchy level.

Copyright © 2015, Juniper Networks, Inc.

Configuring the Continuity Check Loss Threshold

You can specify the number of continuity check messages that can be lost before marking the MEP as down. The default value is three (PDUs).

To configure the loss threshold, include the **loss-threshold** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]** hierarchy level.

Continuity Measurement

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.

Related Documentation

- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- Managing Continuity Measurement Statistics on page 514

- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

Configuring a Maintenance Endpoint

To configure the maintenance endpoint, include the **mep** ***mep-id*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level.

Enabling Maintenance Endpoint Automatic Discovery

You can enable the MEP to accept continuity check messages from all remote MEPs of the same maintenance association.

To configure automatic discovery, include the **auto-discovery** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Configuring the Maintenance Endpoint Direction

You can specify the direction in which CFM packets are transmitted for the MEP.

Direction up continuity check messages (CCMs) are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on this MEP.

Direction down CCMs are transmitted only out of the interface configured on this 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.

To configure the MEP direction, include the **direction** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Copyright © 2015, Juniper Networks, Inc.

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 our 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.

NOTE: If you attempt to disable the control word by configuring 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 for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs, the existing CFM sessions are dropped. To prevent this problem, you must first deactivate the Layer 2 circuit, disable the control word, and reactivate the Layer 2 circuit on both the MEPs of a CFM session.

Configuring the Maintenance Endpoint Interface

You must specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface.

On MX Series routers, you can enable the MEP on a specific VLAN of a trunk interface.

To configure the interface, include the **interface *interface-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

MEP Interface Configuration

This example shows the MEP interface configuration statements:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  domain-name maintenance-association ma-name]
mep mep-id {
  direction (up | down);
  interface (ge | xe)-(fpc/pic/port | fpc/pic/port.domain | fpc/pic/port.domain vlan vlan-id);
  auto-discovery;
  priority number;
}
```

Configuring the Maintenance Endpoint Priority

You can specify the IEEE 802.1 priority bits that are used by continuity check and link trace messages.

 Copyright © 2015, Juniper Networks, Inc.

To configure the priority, include the **priority** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Configuring the Maintenance Endpoint Lowest Priority Defect

You can specify the lowest priority defect that is allowed to generate a fault alarm. This configuration determines whether to generate a fault alarm whenever it detects a defect. This configuration is done at the MEP level.

To configure the lowest priority defect, include the **lowest-priority-defect *options*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Table 30 on page 373 describes the available lowest priority defect options.

Table 30: Lowest Priority Defect Options

Option	Description
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 defect.
rem-err-xcon	Allows only not receiving CCM, erroneous CCM, and cross-connect CCM defects.
xcon	Allows only cross-connect CCM defects.

The following configuration example shows **mac-rem-err-xcon** as the lowest priority defect:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md6 {
        level 6;
        maintenance-association ma6 {
          mep 200 {
            interface ge-5/0/0.0;
            direction down;
            lowest-priority-defect mac-rem-err-xcon;
          }
        }
      }
    }
  }
}
```

Copyright © 2015, Juniper Networks, Inc.

Configuring a Remote Maintenance Endpoint

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

To configure the remote MEP, include the **remote-mep** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Configuring a Remote Maintenance Endpoint Action Profile

You can specify the name of the action profile to use for the remote MEP.

To configure the action profile, include 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 *mep-id*]** hierarchy level. The profile must already be defined at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Configuring Maintenance Endpoint Service Protection

You can enable service protection for a VPWS (Virtual Private Wire Service) 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 MEP **interface** statement within the maintenance association and specify the path as working or protect.

To configure the MEP interface, include the **interface** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*** hierarchy level. On the **interface** statement, specify the path as (**working | protect**). The direction must also be configured as direction down for both sessions.

NOTE: If the path is not specified, the session monitors the active path.

Table 31 on page 374 describes the available service protection options.

Table 31: Service Protection Options

Option	Description
working	Specifies the working path.
protect	Specifies the protect path.

 Copyright © 2015, Juniper Networks, Inc.

The following configuration example shows service protection is enabled for the VPWS service. The CCM session is configured for the working path and references the CCM session configured for the protect path in the **protect-maintenance-association** statement. The APS profile is configured and associated with the maintenance-association for the working path:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain vpws-service-1 {
        name-format none;
        level 5;
        maintenance-association W {
          short-name-format character-string;
            protect-maintenance-association P {
            aps-profile aps-profile-1;
          }
          continuity-check {
            interval 1s;
          }
          mep 1 {
            interface ge-1/3/5.0 working;
            direction down;
            auto-discovery;
          }
        }
        maintenance-association P {
          short-name-format character-string;
          continuity-check {
            interval 1s;
          }
          mep 1 {
            interface ge-1/3/5.0 protect;
            direction down;
            auto-discovery;
          }
        }
      }
    }
  }
}
```

Related Documentation

- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379

Copyright © 2015, Juniper Networks, Inc.

- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring a Connectivity Fault Management Action Profile

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

To configure the action profile name, include the **action-profile** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

- Configuring the Action of a CFM Action Profile on page 376
- Configuring the Default Actions of a CFM Action Profile on page 376
- Configuring a CFM Action Profile Event on page 377

Configuring the Action of a CFM Action Profile

You can configure the action to be taken when any of the configured events occur.

To configure the action profile's action, include the **action** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name*]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile bring-down {
      event {
        interface-status-tlv down;
      }
      action {
        interface-down;
      }
    }
  }
}
```

Configuring the Default Actions of a CFM Action Profile

You can configure the default actions to be taken when connectivity to a remote MEP fails.

To enable the **interface-down** as the default action for an action profile, include the **interface-down** statement at the **[edit protocols oam ethernet**

 Copyright © 2015, Juniper Networks, Inc.

connectivity-fault-management action-profile *profile-name* default-actions] hierarchy level.

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile bring-down {
          default-actions {
            interface-down;
          }
        }
        maintenance-domain md1 {
          level 0;
            maintenance-association ma1 {
            continuity-check {
              interval 100 ms;
            }
            mep 4001 {
              interface ge-4/1/0;
              direction down;
              remote-mep 1 {
                action-profile bring-down;
              }
            }
          }
        }
      }
    }
  }
}
```

NOTE: Associating an **action-profile** with the action of **interface-down** 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.

Configuring a CFM Action Profile Event

You can configure one or more events under the action profile, the occurrence of which triggers the corresponding action to be taken.

To configure the interface-status-tlv lower-layer-down event, include the **interface-status-tlv lower-layer-down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

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 the port-status-tlv blocked event, include the **port-status-tlv blocked** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure the adjacency-loss event, include the **adjacency-loss** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure an RDI event to bring an interface down on reception of an RDI bit from a MEP, include the **rdi** statement at the **(edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile bring-down {
      event {
        adjacency-loss;
        interface-status-tlv (down | lower-layer-down);
        port-status-tlv blocked;
        rdi;
      }
      action {
        interface-down;
      }
      clear-action {
        interface-down peer-interface ;
      }
    }
  }
}
```

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.

Related Documentation

- *event (CFM)*
- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Creating a Maintenance Association on page 367
- Configuring a Maintenance Endpoint on page 371
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 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 for Routing Devices*.

Configuring the Linktrace Path Age Timer

If no response to a **linktrace** request is received, the request and response entries are deleted after the age timer expires. To configure the linktrace age timer, use the **linktrace** statement with the **age** ***time*** option at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level. The age is configured in minutes or seconds.

Configuring the Linktrace Database Size

Configure the number of linktrace reply entries to be stored per linktrace request. To configure the linktrace database size, use the **linktrace** statement with the **path-database-size** ***path-database-size*** option at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Display the linktrace database using the **show oam ethernet connectivity-fault-management path-database** command.

Related Documentation

- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 25 on page 380.

Figure 25: Scope of the E-LMI Protocol

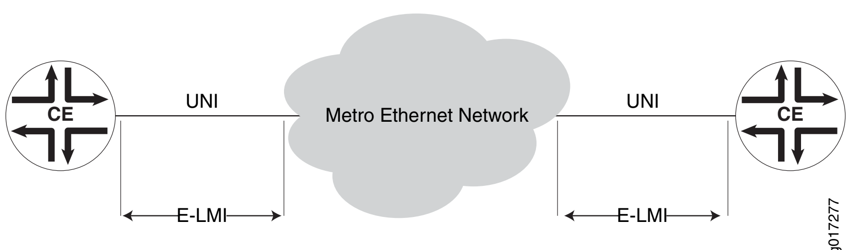

The E-LMI implementation on MX Series routers includes only the PE side of the E-LMI protocol.

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)

 Copyright © 2015, Juniper Networks, Inc.

- 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.**

Configuring the Ethernet Local Management Interface

To configure E-LMI, perform the following steps:

- Configuring an OAM Protocol (CFM) on page 382
- Assigning the OAM Protocol to an EVC on page 382
- Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC on page 382

Configuring an OAM Protocol (CFM)

For information on configuring the OAM protocol (CFM), see "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 360.

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;
```

 Copyright © 2015, Juniper Networks, Inc.

```
            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

Figure 26 on page 383 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 26: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM

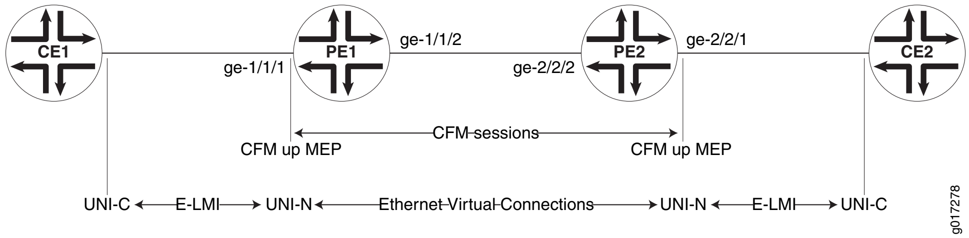

Configuring PE1

```
[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;
                  }
               }
```

 Copyright © 2015, Juniper Networks, Inc.

```
            }
          }
          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;
            }
         }
      }
   }
}
```

 Copyright © 2015, Juniper Networks, Inc.

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 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 32 on page 388 shows the TLV format and indicates if it is required or optional.

Table 32: 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 33 on page 388 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 33: 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

 Copyright © 2015, Juniper Networks, Inc.

Table 33: Type Field Values for Various TLVs for CFM PDUs *(continued)*

TLV or Organization	Type Field
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):

Copyright © 2015, Juniper Networks, Inc.

    - 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.

 Copyright © 2015, Juniper Networks, Inc.

For configuration information, see the following sections:

- Port Status TLV on page 391
- Interface Status TLV on page 393

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 35 on page 391. The format of this TLV is shown in Table 34 on page 391.

Any change in the Port Status TLVs value triggers one extra transmission of that bridge ports MEP CCMs.

Table 34: Port Status TLV Format

Parameter	Octet (Sequence)
Type = 2	1
Length	2–3
Value (See Table 35 on page 391)	4

Table 35: 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.

Copyright © 2015, Juniper Networks, Inc.

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 35 on page 391, 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 36 on page 393. The enumerated values are shown in Table 37 on page 393.

Table 36: Interface Status TLV Format

Parameter	Octet (Sequence)
Type = 4	1
Length	2–3
Value (See Table 37 on page 393)	4

Table 37: Interface Status TLV Values

Mnemonic	Interface Status	Value
isUp	up	1

Table 37: Interface Status TLV Values *(continued)*

Mnemonic	Interface Status	Value
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.

 Copyright © 2015, Juniper Networks, Inc.

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 36 on page 393, 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
```

Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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;
        }
      }
    }
  }
```

 Copyright © 2015, Juniper Networks, Inc.

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 640
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

Figure 27: CET inter-op Dual Homed Topology

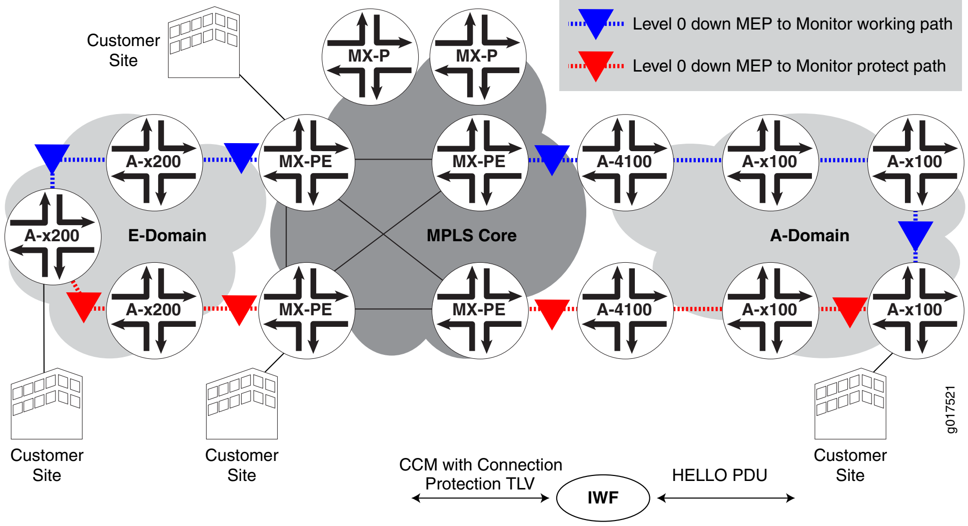

Figure 27 on page 401 describes the dual homed topology on MX Series PE routers connected to the A-domain. When an A-domain device triggers a switchover, it starts switching the service traffic to the new active path. This change is communicated in the HELLO protocol data units (PDUs) sent by that A-domain 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.

Copyright © 2015, Juniper Networks, Inc.

Figure 28: CET inter-op Dual Attached Topology

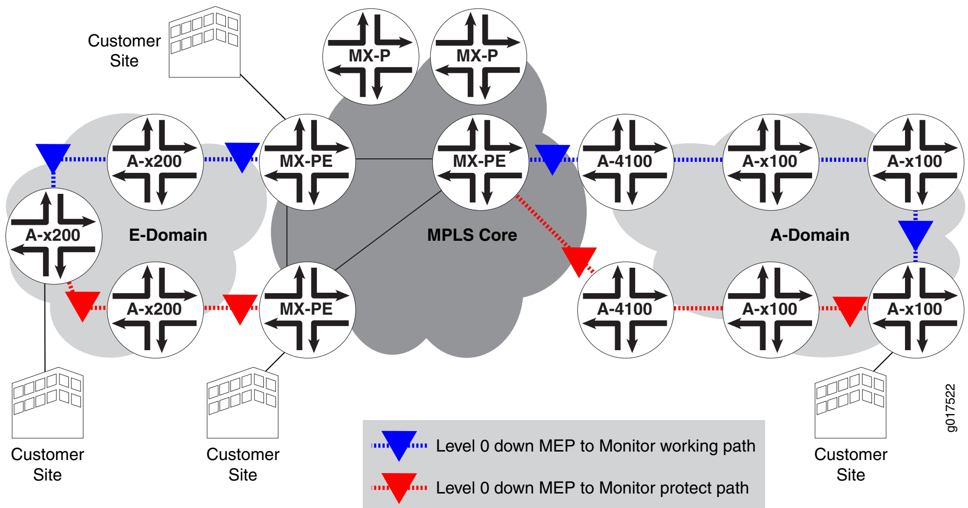

Figure 28 on page 402 describes the dual attached topology on MX Series PE routers connected to the A-domain. The MAC flush mechanism used in this case is also the same as the one used for the A-domain in the dual homed scenario (Figure 1). However in this case both the CFM sessions are hosted by only 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*

 Copyright © 2015, Juniper Networks, Inc.

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 403
- Overview and Topology on page 403
- Configuration on page 404

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 29 on page 404

Figure 29: Topology of CET network

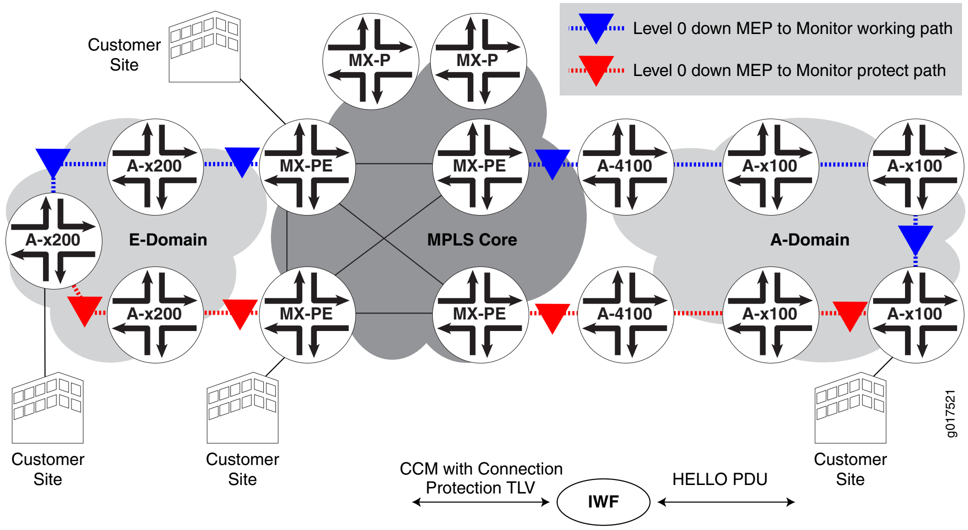

The following definitions describe the meaning of the device abbreviation and terms used in Figure 29 on page 404.

- 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 */
```


Copyright © 2015, Juniper Networks, Inc.

```
        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 400

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.

Copyright © 2015, Juniper Networks, Inc.

- 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 30: Layer 2 VPN Topology

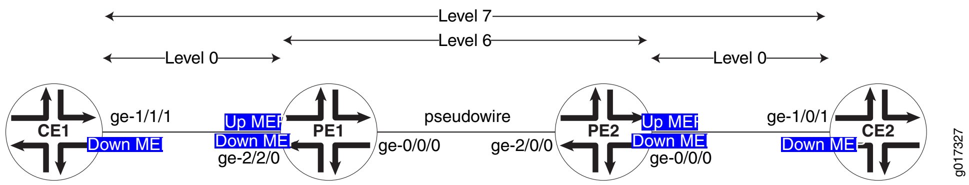

To monitor the L2VPN circuit, a CFM up MEP (Level 6 in Figure 30 on page 406) can be configured on the CE-facing logical interfaces of provider edge routers PE1 and PE2. To monitor the CE-PE attachment circuit, a CFM down MEP can be configured on the customer logical interfaces of CE1-PE1 and CE2-PE2 (Level 0 in Figure 30 on page 406).

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;
          }
        }
      }
    }
  }
  }
}
```

Related Documentation

- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363

Copyright © 2015, Juniper Networks, Inc.

- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring Rate Limiting of Ethernet OAM Messages on page 407
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Rate Limiting of Ethernet OAM Messages

M Series, M320 with Enhanced III FPC, M120, M7i and 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 *name* level *number* maintenance-association *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;
```

Copyright © 2015, Juniper Networks, Inc.

```
  }
}
```

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 **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 **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 **connectivity-fault-management maintenance-domain *md* maintenance-association *ma*** 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 **connectivity-fault-management maintenance-domain *md* maintenance-association *ma*** 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;
        }
      }
    }
  }
```

 Copyright © 2015, Juniper Networks, Inc.

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

- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Creating the Maintenance Domain on page 363
- Configuring Maintenance Intermediate Points on page 365
- Creating a Maintenance Association on page 367
- Continuity Check Protocol on page 368
- Configuring a Maintenance Endpoint on page 371
- Configuring a Connectivity Fault Management Action Profile on page 376
- Configuring Linktrace Protocol in CFM on page 379
- Configuring Ethernet Local Management Interface on page 380
- Configuring Port Status TLV and Interface Status TLV on page 388
- Configuring MAC Flush Message Processing in CET Mode on page 400
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 405
- *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

Copyright © 2015, Juniper Networks, Inc.

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: **Configuring ISSU with CFM (802.1ag) is supported only between two MX 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 360, and you should further observe the additional requirements described in this topic.

Table 38 on page 410 shows the Loss Threshold TLV format.

Table 38: 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)

 Copyright © 2015, Juniper Networks, Inc.

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
```

Copyright © 2015, Juniper Networks, Inc.

```
    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
```

Related Documentation

- Example: Configuring Ethernet CFM over VPLS on page 419
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360

 Copyright © 2015, Juniper Networks, Inc.

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 413.
- 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 415.
- 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 416.
- 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 417.

Related Documentation

- Configuring Faster Protection Switching for Point-to-Point Network Topologies on page 413
- Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies on page 415
- Configuring a Primary VLAN ID for Increased Flexibility on page 416
- Configuring a Remote Maintenance Association to Accept a Different ID on page 417

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
                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 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- Configuring a Connectivity Fault Management Action Profile on page 376

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.

Copyright © 2015, Juniper Networks, Inc.

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 400
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- **connectivity-fault-management on page 640**
- Configuring a Connectivity Fault Management Action Profile on page 376

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;
```


Copyright © 2015, Juniper Networks, Inc.

```
        }
        mep 2 {
            interface ge-2/2/0.0;
            direction up;
            auto-discovery;
        }
    }
}
```

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- *connection-protection-tlv*
- Creating a Maintenance Association on page 367
- **connectivity-fault-management on page 640**

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 {
            interface 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 360
- Creating a Maintenance Association on page 367
- **connectivity-fault-management on page 640**
- *connection-protection-tlv*

Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers

The IEEE 802.1ag provides a specification for Ethernet connectivity fault management (CFM). The Ethernet network may be comprised of one or more service instances. A service instance could be a VLAN, or a concatenation of VLANs. The goal of CFM is to provide a mechanism to monitor, locate, and isolate faulty links. Ethernet 802.1ag is supported on numerous Juniper Networks routers and switches. This topic describes configuration support for Ethernet OAM 802.1ag features on the PTX Series Packet Transport Routers.

Supported features include:

- Maintenance domain (**maintenance-domain *domain-name***) and maintenance levels (**level *number***).
- Maintenance association (**maintenance-association *ma-name***), including name formats (**name-format** and **short-name-format** for **vlan** and **2octet**), loss threshold (**loss-threshold *number***), and hold interval (**hold-interval *minutes***).
- maintenance association endpoint (MEP) functions, including Maintenance Endpoint ID (**mep *mep-id***), direction down (**direction down**), and autodiscovery (**auto-discovery**).
- Link trace for down MEPs (**link-down**).
- action profile (**action-profile *profile-name***)
- Loopback message generation and reply for down MEPs.

Features that are not supported include:

- Up MEP configuration.
- maintenance association intermediate point (MIP) configuration.

To configure flexible Ethernet services encapsulation on PTX Series Packet Transport Routers, include the **oam** statement at the **[edit protocols]** hierarchy level. For example:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md1 {
        level 0;
        maintenance-association ma1 {
          continuity-check {
            interval 100ms;
          }
```


Copyright © 2015, Juniper Networks, Inc.

```
            mep 1 {
              interface et-0/1/1;
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
```

Related Documentation

- Configuring Continuity Check Messages for Better Scalability on page 413
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360

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 31 on page 419. 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 31: Ethernet OAM with VPLS

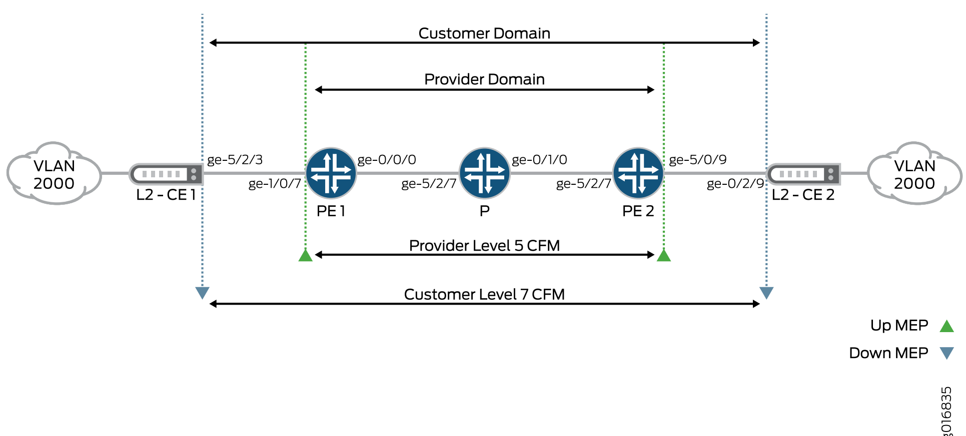

NOTE: The logical interfaces in a VPLS routing instance might have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. Normalization supports automatic mapping of VLANs and performs operations on VLAN tags to achieve the desired 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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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;
      }
```

 Copyright © 2015, Juniper Networks, Inc.

```
        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;
      }
    }
```

 Copyright © 2015, Juniper Networks, Inc.

```
   }
}

[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 OAM Feature Guide for MX Series Routers*
- *Ethernet Operations, Administration, and Maintenance*
- *Ethernet OAM Connectivity Fault Management*
- *Example: Configuring Ethernet CFM on Bridge Connections*
- *Example: Configuring Ethernet CFM on Physical Interfaces*

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 29

Configuring IEEE 802.3ah OAM Link-Fault Management

IEEE 802.3ah OAM Link-Fault Management Overview

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

Copyright © 2015, Juniper Networks, Inc.

technology, as well as being backward-compatible with existing Ethernet technology. Junos OS supports IEEE 802.3ah link-fault management.

The features of link-fault management are:

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

The following features are not supported:

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

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

Related Documentation

- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444

Copyright © 2015, Juniper Networks, Inc.

- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring IEEE 802.3ah OAM Link-Fault Management

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

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

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

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

Related Documentation

- **event-thresholds on page 679**
- **action-profile on page 606**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442

Copyright © 2015, Juniper Networks, Inc.

- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

On PTX Series Packet Transport Routers, Junos OS Release 12.1 supports 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;
        }
      }
    }
  }
}
```

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Configuring Link Discovery on page 432
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437

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 733**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 731**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 798**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the OAM PDU Threshold

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

Copyright © 2015, Juniper Networks, Inc.

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 799
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Threshold Values for Local Fault Events on an Interface

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

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

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

 Copyright © 2015, Juniper Networks, Inc.

```
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
```

Related Documentation

- **event-thresholds on page 679**
- **frame-error on page 695**
- **frame-period on page 696**
- **frame-period-summary on page 697**
- **symbol-period on page 866**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Disabling the Sending of Link Event TLVs

You can disable the sending of link event TLVs.

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

```
[edit protocol oam ethernet link-fault-management interface interface-name
  negotiation-options]
```

no-allow-link-events;

Related Documentation

- **no-allow-link-events on page 775**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Detecting Remote Faults

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

- Critical Event
- Dying Gasp
- Link Fault

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

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

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring an OAM Action Profile

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

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

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

```
  }
}
```

Related Documentation

- **action-profile on page 607**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

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

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

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

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

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

 Copyright © 2015, Juniper Networks, Inc.

```
}
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 604
- syslog on page 866
- link-down on page 731
- send-critical-event on page 843
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444

Copyright © 2015, Juniper Networks, Inc.

- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Monitoring the Loss of Link Adjacency

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

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

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

Related Documentation

- **link-adjacency-loss on page 730**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Monitoring Protocol Status

The CCC-DOWN flag is associated with a circuit cross-connect (CCC) connection, Layer 2 circuit, and Layer 2 VPN, which send the CCC-DOWN status to the kernel. The

 Copyright © 2015, Juniper Networks, Inc.

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:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
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 821**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 732
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445

- *Ethernet Interfaces Feature Guide for Routing Devices*

Applying an Action Profile

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

 Copyright © 2015, Juniper Networks, Inc.

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 619
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Setting a Remote Interface into Loopback Mode

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

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

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

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

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

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

Related Documentation

- **remote-loopback on page 835**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Enabling Remote Loopback Support on the Local Interface on page 444
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling Remote Loopback Support on the Local Interface

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

 Copyright © 2015, Juniper Networks, Inc.

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 619**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 445
- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring IEEE 802.3ah OAM Support on an Interface

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

```
[edit]
protocols {
   oam {
      ethernet {
         link-fault-management {
            interface xe-0/0/0 {
```

Copyright © 2015, Juniper Networks, Inc.

```
                    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 733**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 427
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 429
- Enabling IEEE 802.3ah OAM Support on page 431
- Configuring Link Discovery on page 432
- Configuring the OAM PDU Interval on page 433
- Configuring the OAM PDU Threshold on page 433
- Configuring Threshold Values for Local Fault Events on an Interface on page 434
- Disabling the Sending of Link Event TLVs on page 435
- Detecting Remote Faults on page 436
- Configuring an OAM Action Profile on page 437
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438
- Monitoring the Loss of Link Adjacency on page 440
- Monitoring Protocol Status on page 440
- Configuring Threshold Values for Fault Events in an Action Profile on page 442
- Applying an Action Profile on page 442
- Setting a Remote Interface into Loopback Mode on page 443
- Enabling Remote Loopback Support on the Local Interface on page 444
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 30

Configuring ITU-T Y.1731 Ethernet Service OAM

Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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 32 on page 449 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 32: Relationship of MEPs, MIPs, and Maintenance Domain Levels

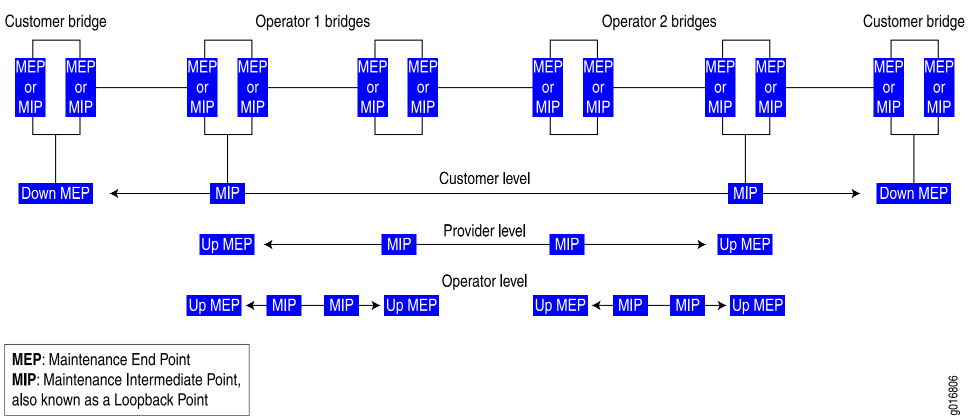

NOTE: Maintenance intermediate points (MIP) are not supported on the ACX Series routers.

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

Copyright © 2015, Juniper Networks, Inc.

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.

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.

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.

 Copyright © 2015, Juniper Networks, Inc.

- 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.

 Copyright © 2015, Juniper Networks, Inc.

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 aggregated Ethernet interfaces or label-switched interface. (LSI) pseudowires.
- 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 488 and "Enabling the Hardware-Assisted Timestamping Option" on page 498.
- 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 488 and "Ensuring That Distributed ppm Is Not Disabled" on page 497.
- 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 499.
- 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 503.
- 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 454
- Example: One-Way Ethernet Frame Delay Measurement on page 515
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- Guidelines for Starting an ETH-DM Session on page 489
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 491
- On-Demand Mode for SLA Measurement on page 456

- Proactive Mode for SLA Measurement on page 457
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 448.

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.

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 Junos OS implementation of Ethernet frame delay measurement (ETH-DM) is fully compliant with the ITU-T Recommendation Y.1731, as described in *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.

Related Documentation

- Managing Continuity Measurement Statistics on page 514
- On-Demand Mode for SLA Measurement on page 456
- Proactive Mode for SLA Measurement on page 457
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 456.
- Proactive mode—In proactive mode, the measurements are triggered by an iterator application. For more information, see "Proactive Mode for SLA Measurement" on page 457.

For more information about frame delay measurement, see "Ethernet Frame Delay Measurements Overview" on page 448. For more information about frame loss measurement, see "Ethernet Frame Loss Measurement Overview" on page 454. 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 457.
- On-Demand Mode for SLA Measurement on page 456.

- *Ethernet Interfaces Feature Guide for Routing Devices*

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 495. 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 448
- Ethernet Frame Loss Measurement Overview on page 454
- Proactive Mode for SLA Measurement on page 457
- Configuring Routers to Support an ETH-DM Session on page 495.
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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.

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 478.

- 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 477.

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 477
- Configuring a Remote MEP with an Iterator Profile on page 486
- Ethernet Frame Delay Measurements Overview on page 448
- Ethernet Frame Loss Measurement Overview on page 454
- Verifying the Configuration of an Iterator Profile on page 478
- Managing Iterator Statistics on page 481
- On-Demand Mode for SLA Measurement on page 456
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

Copyright © 2015, Juniper Networks, Inc.

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 522
- **show oam ethernet fnp interface on page 1402**
- **show oam ethernet fnp status on page 1405**
- **show oam ethernet fnp messages on page 1403**
- **connectivity-fault-management on page 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360
- *Ethernet Interfaces Feature Guide for Routing Devices*

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

Copyright © 2015, Juniper Networks, Inc.

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).

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.

Related Documentation

- Transmission of ETH-SLM Messages on page 464
- Format of ETH-SLM Messages on page 462
- Guidelines for Configuring ETH-SLM on page 466
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Starting an On-Demand ETH-SLM Session on page 471
- Troubleshooting Failures with ETH-SLM on page 476
- *Ethernet Interfaces Feature Guide for Routing Devices*

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:

 Copyright © 2015, Juniper Networks, Inc.

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 459
- Transmission of ETH-SLM Messages on page 464
- Format of ETH-SLM Messages on page 462
- Guidelines for Configuring ETH-SLM on page 466
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Starting an On-Demand ETH-SLM Session on page 471
- Troubleshooting Failures with ETH-SLM on page 476

- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

 Copyright © 2015, Juniper Networks, Inc.

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 459
- Transmission of ETH-SLM Messages on page 464
- Guidelines for Configuring ETH-SLM on page 466
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Starting an On-Demand ETH-SLM Session on page 471

- Troubleshooting Failures with ETH-SLM on page 476
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

 Copyright © 2015, Juniper Networks, Inc.

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 459
- Format of ETH-SLM Messages on page 462

- Guidelines for Configuring ETH-SLM on page 466
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Starting an On-Demand ETH-SLM Session on page 471
- Troubleshooting Failures with ETH-SLM on page 476
- *Ethernet Interfaces Feature Guide for Routing Devices*

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

 Copyright © 2015, Juniper Networks, Inc.

of the downstream MEP must have been configured on the MEP transmitting an alarm indication signal (AIS).

- ETH-SLM is not supported on aggregated Ethernet (ae) interfaces and on downstream MEPs on label-switched interfaces (LSIs).
- 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 459
- Transmission of ETH-SLM Messages on page 464
- Format of ETH-SLM Messages on page 462
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Starting an On-Demand ETH-SLM Session on page 471
- Troubleshooting Failures with ETH-SLM on page 476
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 468
- Configuring an Iterator Profile for ETH-SLM on page 469
- Associating the Iterator Profile with MEPs for ETH-SLM on page 470

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;
             }
           }
         }
       }
     }
   }
   ```

 Copyright © 2015, Juniper Networks, Inc.

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.

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*;

Copyright © 2015, Juniper Networks, Inc.

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*;

Copyright © 2015, Juniper Networks, Inc.

iteration-count *count-value*;
priority *priority-value*;
}

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 459
- Transmission of ETH-SLM Messages on page 464
- Format of ETH-SLM Messages on page 462
- Guidelines for Configuring ETH-SLM on page 466
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting an On-Demand ETH-SLM Session on page 471
- Troubleshooting Failures with ETH-SLM on page 476
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 459
- Transmission of ETH-SLM Messages on page 464

Copyright © 2015, Juniper Networks, Inc.

- Format of ETH-SLM Messages on page 462
- Guidelines for Configuring ETH-SLM on page 466
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Troubleshooting Failures with ETH-SLM on page 476
- *Ethernet Interfaces Feature Guide for Routing Devices*

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

 Copyright © 2015, Juniper Networks, Inc.

this command and the descriptions of the output fields, see **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**.

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**.

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**.

 Copyright © 2015, Juniper Networks, Inc.

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 932
- show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 1398
- show oam ethernet connectivity-fault-management interfaces on page 1349 (**detail** | **extensive**)
- show oam ethernet connectivity-fault-management mep-statistics on page 1377
- show oam ethernet connectivity-fault-management mep-database on page 1366
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 459
- Transmission of ETH-SLM Messages on page 464
- Format of ETH-SLM Messages on page 462
- Guidelines for Configuring ETH-SLM on page 466
- Scenarios for Configuration of ETH-SLM on page 460
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 472
- Starting a Proactive ETH-SLM Session on page 467
- Starting an On-Demand ETH-SLM Session on page 471
- *Ethernet Interfaces Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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).

Copyright © 2015, Juniper Networks, Inc.

[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
user@host# **set calculation-weight delay-variation** ***delay-variation-value***

9. 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**

10. 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);
    calculation-weight {
    delay *delay-weight*;
    delay-variation *delay-variation-weight*;
    }

Related Documentation

- Proactive Mode for SLA Measurement on page 457
- Configuring a Remote MEP with an Iterator Profile on page 486
- Verifying the Configuration of an Iterator Profile on page 478
- Managing Iterator Statistics on page 481
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

 Copyright © 2015, Juniper Networks, Inc.

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 477 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 477 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
```

Copyright © 2015, Juniper Networks, Inc.

```
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 486 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**

 Copyright © 2015, Juniper Networks, Inc.

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 457
- Configuring an Iterator Profile on page 477
- Configuring a Remote MEP with an Iterator Profile on page 486
- Managing Iterator Statistics on page 481

- *Ethernet Interfaces Feature Guide for Routing Devices*

Managing Iterator Statistics

- Displaying Iterator Statistics on page 481
- Clearing Iterator Statistics on page 485

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 39: 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.
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.

Copyright © 2015, Juniper Networks, Inc.

Table 39: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields *(continued)*

Output Field Name	Output Field Description
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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 40: 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.
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

 Copyright © 2015, Juniper Networks, Inc.

Table 40: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields *(continued)*

Output Field Name	Output Field Description
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.

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**

Related Documentation

- Configuring an Iterator Profile on page 477
- Configuring a Remote MEP with an Iterator Profile on page 486
- Verifying the Configuration of an Iterator Profile on page 478
- Proactive Mode for SLA Measurement on page 457

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.

Copyright © 2015, Juniper Networks, Inc.

[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 457
- Configuring an Iterator Profile on page 477
- Verifying the Configuration of an Iterator Profile on page 478
- Managing Iterator Statistics on page 481
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 488.
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 477.

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 488
- Configuring an Iterator Profile on page 477
- Verifying the Configuration of an Iterator Profile on page 478
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 488
- Configuration Options for ETH-DM on page 489

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 Maintenance Endpoint" on page 371.
- 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 367.
- On both routers, the MA must be associated with the same maintenance domain (MD) name. For more information, see "Creating the Maintenance Domain" on page 363.
- 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

Copyright © 2015, Juniper Networks, Inc.

Forwarding Engine of both routers. For more information about **ppm**, see the *Junos OS Routing Protocols Library for Routing Devices*.

- 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 380.

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 448
- Configuring Routers to Support an ETH-DM Session on page 495

- *Ethernet Interfaces Feature Guide for Routing Devices*

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 489
- ETH-DM Session Parameters on page 490
- Restrictions for an ETH-DM Session on page 491

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 the Maintenance Domain" on page 363.

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 488 and "Configuring Routers to Support an ETH-DM Session" on page 495.

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 the Maintenance Domain" on page 363.
  - CFM maintenance association—Name of an existing maintenance association (MA) within the maintenance domain. For more information, see "Creating a Maintenance Association" on page 367.
  - 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 Maintenance Endpoint" on page 371.
- 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.

 Copyright © 2015, Juniper Networks, Inc.

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 448
- Starting an ETH-DM Session on page 499
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 491

- **monitor ethernet delay-measurement** operational command
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 448.
- 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 448.

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 41 on page 492 describes the ETH-DM statistics calculated in an ETH-DM session.

Table 41: 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.

 Copyright © 2015, Juniper Networks, Inc.

Table 41: ETH-DM Statistics *(continued)*

Field Name	Field Description
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.

†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 42 on page 494 describes the ETH-DM frame counts collected in an ETH-DM session.

Table 42: 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.
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*>

 Copyright © 2015, Juniper Networks, Inc.

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 42 on page 494, 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 42 on page 494, 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 448
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 503
- Example: One-Way Ethernet Frame Delay Measurement on page 515
- **clear oam ethernet connectivity-fault-management statistics** command
- **show oam ethernet connectivity-fault-management mep-statistics on page 1377** command
- **show oam ethernet connectivity-fault-management delay-statistics on page 1341** command
- **show oam ethernet connectivity-fault-management interfaces on page 1349** (**detail** | **extensive**) command
- **show oam ethernet connectivity-fault-management mep-database on page 1366** command
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Routers to Support an ETH-DM Session

Copyright © 2015, Juniper Networks, Inc.

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;
               }
            }
         }
      }
   }
}
```

 Copyright © 2015, Juniper Networks, Inc.

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
      ```

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;
      }
```

 Copyright © 2015, Juniper Networks, Inc.

```
        }
      }
    }
```

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;
}
```

Related Documentation

- On-Demand Mode for SLA Measurement on page 456
- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- *Ethernet Interfaces Feature Guide for Routing Devices*
- Ethernet Frame Delay Measurements Overview on page 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- *Ethernet Interfaces Feature Guide for Routing Devices*

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.

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 ---
```

 Copyright © 2015, Juniper Networks, Inc.

```
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.

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

- Ethernet Frame Delay Measurements Overview on page 448
- Guidelines for Starting an ETH-DM Session on page 489
- **monitor ethernet delay-measurement** command
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 491
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 503
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 936

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
```

 Copyright © 2015, Juniper Networks, Inc.

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 936

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
  ```

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**

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**

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.

 Copyright © 2015, Juniper Networks, Inc.

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**

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

- *clear oam ethernet connectivity-fault-management statistics*
- show oam ethernet connectivity-fault-management delay-statistics on page 1341
- show oam ethernet connectivity-fault-management interfaces on page 1349
- show oam ethernet connectivity-fault-management mep-statistics on page 1377
- show oam ethernet connectivity-fault-management mep-database on page 1366
- *Ethernet Interfaces Feature Guide for Routing Devices*

Displaying ETH-DM Statistics Only

Purpose Display ETH-DM statistics.

Copyright © 2015, Juniper Networks, Inc.

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 1341

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 1377

Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

 Copyright © 2015, Juniper Networks, Inc.

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 1366

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 1349**

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 503

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]
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- *Ethernet Interfaces Feature Guide for Routing Devices*

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;
  }
}
```

Copyright © 2015, Juniper Networks, Inc.

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;
     }
     ```

 Copyright © 2015, Juniper Networks, Inc.

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
      ```

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- *Ethernet Interfaces Feature Guide for Routing Devices*

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

 Copyright © 2015, Juniper Networks, Inc.

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 936
- *Junos OS Operational Mode Commands*

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:

Copyright © 2015, Juniper Networks, Inc.

`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**

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**

Related Documentation

- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 503

Managing Continuity Measurement Statistics

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**

 Copyright © 2015, Juniper Networks, Inc.

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***

Related Documentation

- clear oam ethernet connectivity-fault-management continuity-measurement on page 928

- *Ethernet Interfaces Feature Guide for Routing Devices*

- show oam ethernet connectivity-fault-management delay-statistics on page 1341

Example: One-Way Ethernet Frame Delay Measurement

Description of the One-Way Frame Delay Measurement Example

This example shows how you can configure two MX Series routers (**MX-PE1** and **MX-PE2**) to support an ETH-DM session between two peer MEPs (MEP **201** and MEP **101**), initiate a one-way ETH-DM session (from MEP **101** to MEP **201**), and then display the ETH-DM statistics and frame counts collected. To increase the accuracy of the ETH-DM statistics, enable optional hardware-assisted timestamping of received ETH-DM frames on the router that contains the receiver MEP.

Routers Used in This Example

To support one-way ETH-DM with optional hardware timestamping of frames on the reception path, the routers used in this example are configured as follows:

- Routers **MX-PE1** and **MX-PE2** are MX Series routers.
- The system clocks of routers **MX-PE1** and **MX-PE2** are closely synchronized.
- On router **MX-PE1**, interface **ge-5/2/9** is an Ethernet port. The traffic load received is heavy.
- On router **MX-PE2**, interface **ge-0/2/5** is an Ethernet port.

Copyright © 2015, Juniper Networks, Inc.

ETH-DM Frame Counts for this Example

Both routers count the number of ETH-DM frames sent and received by the peer MEPs in the session and store the frame counts in the CFM databases as follows:

- At router **MX-PE2**, which contains the initiator MEP **101**, the CFM database stores the ETH-DM frame counts for a one-way ETH-DM initiator (the count of 1DM frames sent).
- At router **MX-PE1**, which contains the receiver MEP **201**, the CFM database stores the ETH-DM frame counts for a one-way ETH-DM receiver (the count of valid 1DM frames received and the count of invalid 1DM frames received).

ETH-DM Statistics for this Example

For a one-way frame delay measurement, only the router that contains the receiver MEP measures and stores frame delay statistics. In this example, ETH-DM statistics collected for the session are available only at router **MX-PE1**.

 Copyright © 2015, Juniper Networks, Inc.

Steps for the One-Way Frame Delay Measurement Example

The following steps describe an example one-way Ethernet frame delay measurement:

1. At router **MX-PE1**, configure MEP **201** as a CFM maintenance association endpoint in CFM maintenance domain **md6** as follows:

   a. Define the maintenance domain **md6** by associating it with maintenance domain level **6** and maintenance association identifier **ma6**.

   b. Configure the maintenance association by specifying continuity protocol options and specifying MEP identifier **201**.

   c. Configure MEP **201** by attaching it to logical interface **ge-5/2/9.0**, which is a single-tag interface on VLAN **512**.

   The following configuration is only a partial example of a complete and functional router configuration:

```
[edit]
interfaces { # Configure a single-tag logical interface on VLAN 512
  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;
        }
        maintenance-domain md6 { # Define MD 'md6' on router MX-PE1
          level 6;
          maintenance-association ma6 { # Configure MA 'ma6' on router MX-PE1
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 201 { # Configure MEP 201 on router MX-PE1
              interface ge-5/2/9.0; # Attach to logical interface on VLAN 512
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

2. At router **MX-PE2**, configure MEP **101** as a CFM maintenance association endpoint in CFM maintenance domain **md6** as follows:

   a. Define the maintenance domain **md6** by associating it with maintenance domain level **6** and maintenance association identifier **ma6**.

   b. Configure the maintenance association by specifying continuity protocol options and specifying MEP identifier **101**.

   c. Configure MEP **101** by attaching it to logical interface **ge-0/2/5.0**, which is a single-tag interface on VLAN **512**.

   The following configuration is only a partial example of a complete and functional configuration for router **MX-PE2**:

```
[edit]
interfaces { # Configure a single-tag logical interface on VLAN 512
  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;
        }
        maintenance-domain md6 { # Define MD 'md6' on router MX-PE2
          level 6;
          maintenance-association ma6 { # Configure MA 'ma6' on router MX-PE2
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 101 { # Configure MEP 101 on router MX-PE2
              interface ge–0/2/5.0; # Attach to logical interface on VLAN 512
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
```

3. (Optional) To increase the accuracy of the ETH-DM statistics, modify the configuration of router **MX-PE1**, which contains the receiver MEP, by enabling hardware-assisted timestamping of **1DM** frames received on the router.

```
[edit protocols]
oam {
```

 Copyright © 2015, Juniper Networks, Inc.

```
    ethernet {
      connectivity-fault-management {
        performance-monitoring {
          hardware-assisted-timestamping;
        }
      }
    }
  }
```

4. At router **MX-PE2**, start a one-way frame delay measurement session from local MEP **101** to remote MEP **201** on router **MX-PE1**:

```
user@MX-PE2> 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
```

5. At router **MX-PE2**, which contains the initiator MEP, only the ETH-DM frame counts are available. Furthermore, the only frame count tallied for the initiator of a one-way frame delay measurement is the count of 1DM frames transmitted.

   ETH-DM frame counts (the number of 1DM, DMM, and DMR frames exchanged during an ETH-DM session) are stored in the CFM database of both the initiator and receiver MEPs. When you display CFM database information, you can also display the ETH-DM frame counts. You can display CFM database information for all interfaces on the router, or you can limit the output to MEPs associated with certain CFM MDs and MAs.

   - To display CFM database information for MEPs specified by enclosing CFM entities, use the **mep-database** form of the **show oam ethernet connectivity-fault-management** command. A CFM database also stores any ETH-DM frame counts.

     In the example configuration for router **MX-PE2**, MEP **101** is the only MEP defined in MA **ma6** within MD **md6**. Therefore, the **show oam ethernet connectivity-fault management mep-database** command output displays CFM database information for MEP **101** only, even though you do not filter the command output by including the **local-mep** or **remote-mep** command options.

```
user@MX-PE2> 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: 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
```

- To display CFM database information for MEPs specified by interface name, use the **interfaces detail** form of the **show oam ethernet connectivity-fault-management** command. A CFM database also stores any ETH-DM frame counts.

  In the example configuration for router **MX-PE2**, MEP **101** is the only MEP assigned to an interface on the router. Therefore, the **show oam ethernet connectivity-fault-management interfaces** (**detail** | **extensive**) command output displays CFM database information for MEP **101** only, even though you do not filter the command output by including the ***ethernet-interface-name*** or **level** ***md-level*** command options.

```
user@MX-PE2>  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          : 0
      1DMs sent                                    : 10
      Valid 1DMs received                          : 0
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

NOTE: **You can use these same commands—show oam ethernet connectivity-fault-management mep-database and show oam ethernet connectivity-fault-management interfaces (detail | extensive)—at router MX-PE1 to display the CFM database information (which includes any ETH-DM frame counts) for receiver MEP 201.**

6. At router **MX-PE1**, which contains the receiver MEP, you can use two different **show oam ethernet connectivity-fault-management** commands to display ETH-DM statistics and ETH-DM frame counts.

   - To display only the delay statistics, use the **delay-statistics** form of the **show oam ethernet connectivity-fault-management** command:

     `user@MX-PE1>` **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
    Worst case one-way delay       : 370 usec
```

   - To display both the ETH-DM statistics and the CFM database information (which includes any ETH-DM frame counts), use the **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command:

     `user@MX-PE1>` **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
```

Related Documentation

- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- Guidelines for Starting an ETH-DM Session on page 489
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 491
- On-Demand Mode for SLA Measurement on page 456
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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 640**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

PART 4

Troubleshooting Information

Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 31

Monitoring and Troubleshooting Ethernet Interfaces

Interface Diagnostics Tools

You can use two diagnostic tools to test the physical layer connections of interfaces: loopback testing and bit error rate test (BERT) testing. Loopback testing enables you to verify the connectivity of a circuit. BERT testing enables you to identify poor signal quality on a circuit. This section contains the following topics:

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

Copyright © 2015, Juniper Networks, Inc.

troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

Thefollowing types of loopback testing are supported by the 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 43 on page 528 shows the loopback modes supported on the various interface types.

Table 43: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	"Configuring Ethernet Loopback Capability" on page 19
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 Channelized E1 IQ and IQE Interfaces*, *Configuring T1 and NxDS0 Interfaces*, *Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)*, *Configuring Channelized STM1 IQ and IQE Interfaces*, and *Configuring Channelized T3 IQ Interfaces*

 Copyright © 2015, Juniper Networks, Inc.

Table 43: 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*
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:

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]**

BERT Testing

A bit error rate test (BERT) allows you to troubleshoot problems by checking the quality of links. You can configure any of the following interfaces to execute a BERT when the interface receives a request to run this test: E1, E3, T1, T3; the channelized DS3, OC3, OC12, and STM1 interfaces; and the channelized DS3 IQ, E1 IQ, and OC12 IQ interfaces.

A BERT test requires a line loop to be in place on either the transmission devices or the far-end router. The local router generates a known bit pattern and sends it out the transmit path. The received pattern is then verified against the sent pattern. The higher the bit error rate of the received pattern, the worse the noise is on the physical circuit. As you move the position of the line loop increasingly downstream toward the far-end router, you can isolate the troubled portion of the link.

Copyright © 2015, Juniper Networks, Inc.

To configure BERT, you must configure the duration of the test, the bit pattern to send on the transmit path, and the error rate to monitor when the inbound pattern is received.

To configure the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream, include the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements, respectively, at the **[edit interfaces *interface-name interface-type*-options]** hierarchy level:

[edit interfaces *interface-name interface-type*-options]
bert-algorithm *algorithm*;
bert-error-rate *rate*;
bert-period *seconds*;

By default, the BERT period is 10 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.

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).

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.

 Copyright © 2015, Juniper Networks, Inc.

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 44 on page 532 shows the BERT capabilities for various interface types.

Table 44: 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

Copyright © 2015, Juniper Networks, Inc.

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.

Example: Configuring Bit Error Rate Testing

Configure a BERT test on a T3 interface. In this example, the run duration lasts for 120 seconds. The configured error rate is 0, which corresponds to a bit error rate of 10^{-0} (1 error per bit). The configured bit pattern of **all-ones-repeating** means that every bit the interface sends is set to a value of 1.

```
[edit interfaces]
t3-1/2/0 {
  t3-options {
    bert algorithm all-ones-repeating;
    bert-error-rate 0;
    bert-period 120;
```

Copyright © 2015, Juniper Networks, Inc.

```
    }
  }
```

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 45 on page 534 provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

Table 45: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
“Monitor Fast Ethernet and Gigabit Ethernet Interfaces” on page 534	
1. Display the Status of Fast Ethernet Interfaces on page 535	**show interfaces terse (fe* | ge*)**
2. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 537	**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 538	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
4. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 541	**monitor interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
5. Fiber-Optic Ethernet Interface Specifications on page 542	

Meaning You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

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 535
2. Display the Status of Gigabit Ethernet Interfaces on page 536
3. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 537

 Copyright © 2015, Juniper Networks, Inc.

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 46 on page 535.

Table 46: 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.

Table 46: Status of Fast Ethernet Interfaces *(continued)*

Physical Interface	Logical Interface	Status Description
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.

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 47 on page 536 for a description of what the output means.

Table 47: 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.

Copyright © 2015, Juniper Networks, Inc.

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

Copyright © 2015, Juniper Networks, Inc.

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 538, 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 48 on page 539 for a description of errors to look for.

Table 48: 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 553 for an explanation of Fast Ethernet and Gigabit Ethernet alarms.

Table 49 on page 540 lists and describes some MAC statistics errors to look for.

Copyright © 2015, Juniper Networks, Inc.

Table 49: 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 50 on page 540 explains the autonegotiation information of the **show interface *interface-name* extensive** command output.

Table 50: 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.

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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 542 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 542 for the optical specifications.

Fiber-Optic Ethernet Interface Specifications

Table 51 on page 542 shows the specifications for fiber-optic interfaces for Juniper Networks routers.

Table 51: 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					

Copyright © 2015, Juniper Networks, Inc.

Table 51: 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

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 52 on page 543 provides links and commands for using loopback testing for Fast Ethernet and Gigabit Ethernet interfaces.

Table 52: 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 544	
1. Create a Loopback on page 544	
a. Create a Physical Loopback for a Fiber-Optic Interface on page 545	Connect the transmit port to the receive port.
b. Create a Loopback Plug for an RJ-45 Ethernet Interface on page 545	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 546	**[edit interfaces *interface-name* (fastether-options | gigether-options)]** **set loopback** **show** **commit**

Table 52: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces *(continued)*

Tasks	Command or Action
2. Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 547	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
3. Configure a Static Address Resolution Protocol Table Entry on page 548	**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 550	**clear interfaces statistics fe-*fpc/pic/port* | ge-*fpc/pic/port***
5. Ping the Fast Ethernet or Gigabit Ethernet Interface on page 550	**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 551	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
"Diagnose a Suspected Circuit Problem" on page 553	Perform Steps 2 through 8 from "Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 544.

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 544
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 547
- Configure a Static Address Resolution Protocol Table Entry on page 548
- Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 550
- Ping the Fast Ethernet or Gigabit Ethernet Interface on page 550
- Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 551

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

 Copyright © 2015, Juniper Networks, Inc.

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 545
2. Create a Loopback Plug for an RJ-45 Ethernet Interface on page 545
3. Configure a Local Loopback on page 546

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.

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 33 on page 546 illustrates how to create a loopback plug for an RJ-45 Ethernet interface.

Copyright © 2015, Juniper Networks, Inc.

Figure 33: RJ-45 Ethernet Loopback Plug

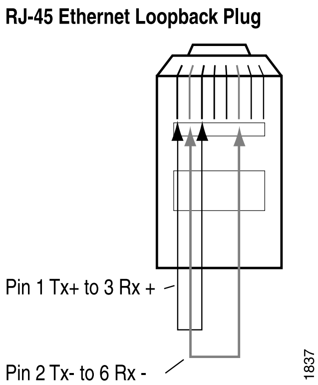

Meaning

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.

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

 Copyright © 2015, Juniper Networks, Inc.

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.**

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 fe-1/3/0
Physical interface: fe-1/3/0, Enabled, Physical link is Up
  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
  Interface flags: 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 : None
  Active defects : None
  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: OK
    Link partner: Full-duplex, Flow control: None
```

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 53 on page 548 presents problem situations and actions for a physical link that is down.

Table 53: 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.

 Copyright © 2015, Juniper Networks, Inc.

NOTE: Remove the static ARP entry at the end of the loop test after you have completed the ping test, checked interface statistics, 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***

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***

   NOTE: The MAC address used should be the same as the physical address of the port being tested because this allows the port to receive the frames when you run the ping test.

4. Verify the configuration:

   user@host# **show**

5. Commit the configuration:

   user@host# **commit**

6. Verify that the static ARP entry is installed:

   user@host# **run show arp no-resolve**

Sample Output

```
user@host> show interfaces ge-7/2/1
Physical interface: ge-7/2/1, 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)
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host#  set arp 10.108.120.2 mac 00:90:69:8d:2c:db
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]
```

Copyright © 2015, Juniper Networks, Inc.

```
user@host# show
arp 10.108.120.2 mac 00:90:69:8d:2c:db;
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host# commit
commit complete
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host# run show arp no-resolve
MAC Address       Address         Interface     Flags
00:90:69:8d:2c:db 10.108.120.2    ge-7/2/1.0    permanent
00:e0:34:bb:8c:40 209.211.135.1   fxp0.0        none
00:a0:a5:28:0c:70 209.211.135.8   fxp0.0        none
00:a0:a5:12:12:c7 209.211.135.10  fxp0.0        none
00:90:ab:3c:68:a0 209.211.135.31  fxp0.0        none
08:00:20:a1:53:15 209.211.135.65  fxp0.0        none
00:a0:cc:66:3e:85 209.211.135.98  fxp0.0        none
Total entries: 7
```

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-7/2/1**. The MAC address used is the same as the physical address of the port being tested because this allows the port to receive the frames when you run the ping test. The port is working as expected if you see that the time to live (TTL) expired; if you do not receive a response to your ping test, it indicates a hardware problem.

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics

Purpose

You must reset the Fast Ethernet and Gigabit Ethernet interface statistics before initiating the ping test. 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-7/2/0
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.

 Copyright © 2015, Juniper Networks, Inc.

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-7/2/1 extensive
Physical interface: ge-7/2/1, Enabled, Physical link is Up
  Interface index: 25, SNMP ifIndex: 32, Generation: 41
  Description: Test
  Link-level type: Ethernet, MTU: 4470, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:4c:17:b1, Hardware address: 00:90:69:4c:17:b1
  Statistics last cleared: 2002-01-07 17:53:19 UTC (2w2d 03:20 ago)
  Traffic statistics:
   Input  bytes  :       3799515503823                    0 bps
   Output bytes  :          7325566425                    0 bps
   Input  packets:          4628009535                    0 pps
   Output packets:            30678225                    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
  Output errors:
    Carrier transitions: 14, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0,
    FIFO errors: 0, HS link CRC errors: 0
  Active alarms : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                 3883579444813       7880356346
    Total packets                   4628009534         30678237
    Unicast packets                 4627879788         29893563
    Broadcast packets                       30              464
    Multicast packets                   129716           784210
    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              4628009244
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                30678237
    Output packet pad count                              856248
    Output packet error count                                 0
    CAM destination filters: 9, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: Ok, Link partner:
Full-duplex,
    Flow control: None
  PFE configuration:
    Destination slot: 7
    CoS transmit queue          Bandwidth          Buffer      Priority   Limit
```

 Copyright © 2015, Juniper Networks, Inc.

```
                               %          bps   %        bytes
    0 best-effort              0            0   0            0      low     none
    1 expedited-forwarding     0            0   0            0      low     none
    2 assured-forwarding       0            0   0            0      low     none
    3 network-control          0            0   0            0      low     none
  Logical interface ge-7/2/1.0 (Index 23) (SNMP ifIndex 48) (Generation 38)
    Description: To Cosine Left 23/1
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 4456, Flags: None, Generation: 85 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.108.120.0/30, Local: 10.108.120.1, Broadcast: 10.108.120.3,

        Generation: 81
    Protocol iso, MTU: 4453, Flags: None, Generation: 86 Route table: 0
```

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 544. 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 54 on page 554 provides links and commands for locating LINK alarm and major counters for Fast Ethernet and Gigabit Ethernet interfaces.

Table 54: 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 554	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
"Fast Ethernet and Gigabit Ethernet Counters" on page 555	

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 543 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.**

Fast Ethernet and Gigabit Ethernet Counters

Problem **Description:** Table 55 on page 556 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.

Solution

Table 55: 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.

 Copyright © 2015, Juniper Networks, Inc.

Table 55: 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.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

PART 5

Configuration Statements and Operational Commands

 Copyright © 2015, Juniper Networks, Inc.

CHAPTER 32

Configuration Statements

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

[edit chassis] Hierarchy Level

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
    sonet {
      device-count number;
    }
  }
    maximum-links {
    }
  channel-group number {
    ethernet {
      device-count number;
    }
    fpc slot-number{
      pic pic-number {
        adaptive-services{
          service-package (layer-2 | layer-3);
        }
        aggregate-ports;
```

 Copyright © 2015, Juniper Networks, Inc.

```
            atm-cell-relay-accumulation;
            atm-l2circuit-mode (aal5 | cell | trunk trunk);
            ce1 {
              e1 link-number {
                channel-group group-number;
                timeslots time-slot-range;
              }
            }
            channelization;
            ct1 {
              t1 link-number {
                channel-group group-number;
                timeslots time-slot-range;
              }
            }
            ct3 {
              port port-number {
                t1 link-number {
                  channel-group group-number;
                  timeslots time-slot-range;
                }
              }
              framing sdh;
            }
            max-queues-per-interface number;
            mlfr-uni-nni-bundles num-intf;
            no-concatenate;
            shdsl {
              pic-mode (1-port-atm | 2-port-atm);
            }
            vtmapping (klm | itu-t);
          }
        }
        fpc slot-number{
        pic pic-number{
          account-layer2-overhead
            egress-policer-overhead bytes;
            ingress-policer-overhead bytes;
            mlfr-uni-nni-bundles-inline number;
            multi-link-layer-2-inline;
          }
        }
      }
```

Related Documentation

- *Interfaces Fundamentals for Routing Devices*

[edit interfaces] Hierarchy Level

The statements at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level can also be configured at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Copyright © 2015, Juniper Networks, Inc.

NOTE: The *accounting-profile* statement is an exception to this rule. The **accounting-profile** statement can be configured at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, but it cannot be configured at the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

```
interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
        (revertive | non-revertive);
        periodic interval;
        system-priority priority;
      }
      link-protection;
      link-speed speed;
      (loopback | no-loopback);
      mc-ae{
        chassis-id chassis-id;
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
      minimum-links number;
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
    shared-scheduler;
    aggregated-sonet-options {
      link-speed speed | mixed;
      minimum-links number;
    }
    atm-options {
      cell-bundle-size cells;
      ilmi;
      linear-red-profiles profile-name {
        high-plp-max-threshold percent;
        low-plp-max-threshold percent;
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;
      }
```

 Copyright © 2015, Juniper Networks, Inc.

```
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    pic-type (atm1 | atm2);
    plp-to-clp;
    promiscuous-mode {
      vpi vpi-identifier;
    }
    scheduler-maps map-name {
      forwarding-class class-name {
        epd-threshold cells plp1 cells;
        linear-red-profile profile-name;
        priority (high | low);
        transmit-weight (cells number | percent number);
      }
      vc-cos-mode (alternate | strict);
    }
    use-null-cw;
      vpi vpi-identifier {
      maximum-vcs maximum-vcs;
      oam-liveness {
        down-count cells;
        up-count cells;
      }
      oam-period (seconds | disable);
      shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
          burst length);
        queue-length number;
      }
    }
  }
  clocking clock-source;
  damping {
    enable;
    half-life seconds;
    max-suppress seconds;
    reuse number;
    suppress number;
  }
  data-input (system | interface interface-name);
  dce;
  serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
      control-signal (assert | de-assert | normal);
      cts (ignore | normal | require);
      dcd (ignore | normal | require);
      dsr (ignore | normal | require);
      dtr signal-handling-option;
```

```
      ignore-all;
      indication (ignore | normal | require);
      rts (assert | de-assert | normal);
      tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
      control-signal (assert | de-assert | normal);
      cts (ignore | normal | require);
      dcd (ignore | normal | require);
      dsr (ignore | normal | require);
      dtr signal-handling-option;
      ignore-all;
      indication (ignore | normal | require);
      rts (assert | de-assert | normal);
      tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
  }
  description text;
  alias alias-name;
  dialer-options {
    pool pool-name <priority priority>;
  }
  disable;
  ds0-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback payload;
    start-end-flag (filler | shared);
  }
  e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (16 | 32);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback (local | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-range;
  }
  e3-options {
```

```
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    framing feet;
    compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
    fcs (16 | 32);
    framing (g.751 | g.832);
    idle-cycle-flag (filler | shared);
    invert-data;
    loopback (local | remote);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
    (unframed | no-unframed);
  }
  encapsulation type;
  es-options {
    backup-interface es-fpc/pic/port;
  }
  fastether-options {
    802.3ad aex;
    (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);
  }
  flexible-vlan-tagging;
  gigether-options {
    802.3ad aex;
    (asynchronous-notification | no-asynchronous-notification);
    (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
      local-interface-offline>;
    auto-reconnect seconds;
    (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);
    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;
            }
          }
        }
      }
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    ima-group-options {
      differential-delay number;
      frame-length (32 | 64 | 128 | 256);
      frame-synchronization {
        alpha number;
        beta number;
        gamma number;
      }
      minimum-links number;
      symmetry (symmetrical-config-and-operation |
        symmetrical-config-asymmetrical-operation);
      test-procedure {
        ima-test-start;
        ima-test-stop;
        interface name;
        pattern number;
        period number;
      }
      transmit-clock (common | independent);
      version (1.0 |1.1);
    }
    ima-link-options group-id group-id;
    interface-set interface-set-name {
      interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
      }
```

 Copyright © 2015, Juniper Networks, Inc.

```
    interface interface-name {
      (unit unit-number);
    }
  }
  isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
  }
  keepalives <down-count number> <interval seconds> <up-count number>;
  link-mode mode;
  lmi {
    lmi-type (ansi | itu | c-lmi);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
  }
  lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
  }
  mac mac-address;
  mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    cisco-interoperability send-lip-remove-link-for-link-reject;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu | c-lmi);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
  }
  modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
  }
```

Copyright © 2015, Juniper Networks, Inc.

```
mtu bytes;
multi-chassis-protection {
  peer a.b.c.d {
    interfaceinterface-name;
    }
  }
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
  interface-type type;
}
no-vpivci-swapping;
otn-options {
  fec (efec | gfec | none);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  pass-thru;
  rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
  transmit-payload-type number;
  trigger (oc-lof | oc-lom | oc-los | oc-wavelength-lock | odu-ais | odu-bbe-th | odu-bdi
    | odu-es-th | odu-lck | odu-oci | odu-sd | odu-ses-th | odu-ttim | odu-uas-th |
    opu-ptm | otu-ais | otu-bbe-th | otu-bdi | otu-es-th | otu-fec-deg | otu-fec-exe |
    otu-iae | otu-sd | otu-ses-th | otu-ttim | otu-uas-th);
  tti;
}
optics-options {
  wavelength nm;
  alarm alarm-name {
    (syslog | link-down);
  }
  warning warning-name {
    (syslog | link-down);
  }
}
partition partition-number oc-slice oc-slice-range interface-type type;
timeslots time-slot-range;
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  no-termination-request;
```

 Copyright © 2015, Juniper Networks, Inc.

```
    pap {
      access-profile name;
      local-name name;
      local-password password;
      compression;
    }
  }
  psn-vcipsn-vci-identifier;
  psn-vpipsn-vpi-identifier;
  receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
  }
  redundancy-options {
    priority sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby;
  }
  satop-options {
    payload-size n;
  }
  schedulers number;
  serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
      control-signal (assert | de-assert | normal);
      cts (ignore | normal | require);
      dcd (ignore | normal | require);
      dsr (ignore | normal | require);
      dtr signal-handling-option;
      ignore-all;
      indication (ignore | normal | require);
      rts (assert | de-assert | normal);
      tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
      control-signal (assert | de-assert | normal);
      cts (ignore | normal | require);
      dcd (ignore | normal | require);
      dsr (ignore | normal | require);
      dtr signal-handling-option;
      ignore-all;
      indication (ignore | normal | require);
      rts (assert | de-assert | normal);
      tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
```

```
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
  }
  services-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    session-limit {
      maximum number;
      rate new-sessions-per-second;
      }
      syslog {
      host hostname {
        facility-override facility-name;
        log-prefix prefix-number;
        services priority-level;
      }
    }
  }
  shdsl-options {
    annex (annex-a | annex-b);
    line-rate line-rate;
    loopback (local | remote);
    snr-margin {
      current margin;
      snext margin;
    }
  }
  sonet-options {
    aggregate asx;
    aps {
      advertise-interval milliseconds;
      annex-b;
      authentication-key key;
      fast-aps-switch;
      force;
      hold-time milliseconds;
      lockout;
      neighbor address;
      paired-group group-name;
      preserve-interface;
      protect-circuit group-name;
      request;
      revert-time seconds;
      switching-mode (bidirectional | unidirectional);
      working-circuit group-name;
    }
    bytes {
      c2 value;
      e1-quiet value;
      f1 value;
      f2 value;
      s1 value;
      z3 value;
```

 Copyright © 2015, Juniper Networks, Inc.

```
      z4 value;
    }
    fcs (16 | 32);
    loopback (local | remote);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    path-trace trace-string;
    (payload-scrambler | no-payload-scrambler);
    rfc-2615;
    trigger {
      defect ignore;
      hold-time up milliseconds down milliseconds;
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
  }
  speed (10m | 100m | 1g | oc3 | oc12 | oc48);
  stacked-vlan-tagging;
  switch-options {
    switch-port port-number {
      (auto-negotiation | no-auto-negotiation);
      speed (10m | 100m | 1g);
      link-mode (full-duplex | half-duplex);
    }
  }
  t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flag (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
  }
  t3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
      value>;
    fcs (16 | 32);
```

```
  (feac-loop-respond | no-feac-loop-respond);
  idle-cycle-flag value;
  (long-buildout | no-long-buildout);
  (loop-timing | no-loop-timing);
  loopback (local | payload | remote);
  (mac | no-mac);
  (payload-scrambler | no-payload-scrambler);
  start-end-flag (filler | shared);
}
traceoptions {
  flag flag <flag-modifier> <disable>;
}
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
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);
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      maximum-contexts number <force>;
      queues [ queue-numbers ];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
  compression-device interface-name;
  copy-tos-to-outer-ip-header;
```

 Copyright © 2015, Juniper Networks, Inc.

```
demux-destination family;
demux-source family;
demux-options {
   underlying-interface interface-name;
}
description text;
alias alias-name;
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;
fragment-threshold bytes;
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 {
  down-count cells;
  up-count cells;
}
oam-period (seconds | disable);
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;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
    pap;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
  dynamic-profile profile-name;
  lcp-max-conf-req number;
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-max-conf-req number;
  ncp-restart-timer milliseconds;
}
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
proxy-arp;
service-domain (inside | outside);
```

 Copyright © 2015, Juniper Networks, Inc.

```
shaping {
   (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
      burst length);
   queue-length number;
}
short-sequence;
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-list [vlan-id vlan-id–vlan-id];
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags-outer tpid.vlan-id inner-list [vlan-id vlan-id–vlan-id];
family family {
   accounting {
      destination-class-usage;
      source-class-usage {
         direction;
      }
   }
   access-concentrator name;
   address address {
      destination address;
   }
   bundle ml-fpc/pic/port | ls-fpc/pic/port);
   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 ];
      }
   }
   ipsec-sa sa-name;
   keep-address-and-control;
```

Copyright © 2015, Juniper Networks, Inc.

```
max-sessions number;
max-sessions-vsa-ignore;
mtu bytes;
multicast-only;
negotiate-address;
no-redirects;
policer {
  arp policer-template-name;
  disable-arp-policer
  input policer-template-name;
  output policer-template-name;
}
primary;
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 {
  direction;
}
service {
  input {
    service-set service-set-name <service-filter filter-name>;
    post-service-filter filter-name;
  }
  output {
    service-set service-set-names <service-filter filter-name>;
  }
}
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
  maximum-seconds>;
targeted-broadcast {
  forward-and-send-to-re;
  forward-only;
}
(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 | preferred-source-address address>;
address address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
  multipoint-destination address {
    epd-threshold cells plp1 cells;
    inverse-arp;
    oam-liveness {
```

 Copyright © 2015, Juniper Networks, Inc.

```
                up-count cells;
                down-count cells;
              }
              oam-period (seconds | disable);
              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;
            }
            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 number;
              track {
                priority-cost seconds;
                priority-hold-time interface-name {
                  bandwidth-threshold bits-per-second {
                    priority;
                  }
                  interface priority;
                }
                route ip-address/mask routing-instance instance-name priority-cost cost;
              }
              virtual-address [ addresses ];
            }
          }
        }
      }
    }
  }
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit logical-systems] Hierarchy Level

The following lists the statements that can be configured at the **[edit logical-systems]** hierarchy level that are also documented in this manual. For more information about logical systems, see the *Logical Systems Feature Guide for Routing Devices*.

```
logical-systems logical-system-name {
  interfaces interface-name {
```

Copyright © 2015, Juniper Networks, Inc.

```
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  bandwidth rate;
  backup-options {
    interface interface-name;
  }
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      port {
        minimum port-number;
        maximum port-number;
      }
      queues [ queue-numbers ];
    }
  }
  compression-device interface-name;
  description text;
  interface {
    l2tp-interface-id name;
    (dedicated | shared);
  }
  dialer-options {
    activation-delay seconds;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    initial-route-check seconds;
    load-threshold number;
    pool pool;
    remote-name remote-callers;
    watch-list {
      [ routes ];
    }
  }
  disable;
  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;
  fragment-threshold bytes;
```

 Copyright © 2015, Juniper Networks, Inc.

```
input-vlan-map {
  inner-tag-protocol-id;
  inner-vlan-id;
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  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 (seconds | disable);
output-vlan-map {
  inner-tag-protocol-id;
  inner-vlan-id;
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
  tag-protocol-id tpid;
  vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
    }
  }
  dynamic-profile profile-name;
  pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
```

```
    }
  }
  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;
  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;
  vlan-id number;
  vlan-id-list [vlan-id vlan-id–vlan-id]
  vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
  vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id–vlan-id]
  vpi vpi-identifier;
  family family {
    accounting {
      destination-class-usage;
      source-class-usage {
        direction;
      }
    }
    bundle interface-name;
    filter {
      group filter-group-number;
      input filter-name;
      input-list {
        [ filter-names ];
      }
      output filter-name;
      output-list {
        [ filter-names ];
      }
    }
    ipsec-sa sa-name;
    keep-address-and-control;
    mtu bytes;
    multicast-only;
    no-redirects;
    policer {
      arp policer-template-name;
```

 Copyright © 2015, Juniper Networks, Inc.

```
    input policer-template-name;
    output policer-template-name;
}
primary;
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 {
    direction;
}
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>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        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;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise–interval seconds;
        authentication-type authentication;
        authentication-key key;
```

Copyright © 2015, Juniper Networks, Inc.

```
                    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 ];
                }
              }
            }
          }
        }
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols connections] Hierarchy Level

The following statements can also be configured at the **[edit logical-systems *logical-system-name* protocols connections]** hierarchy level.

```
interface-switch connection-name {
  interface interface-name.unit-number;
  interface interface-name.unit-number;
}
```

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols dot1x] Hierarchy Level

```
dot1x {
  authenticator
    authentication-profile-name access-profile-name;
    interface interface-ids {
      maximum-requests integer;
      retries integer;
      quiet-period seconds;
      transmit-period seconds;
      reauthentication (disable | interval seconds);
      server-timeout seconds;
      supplicant (single);
      supplicant-timeout seconds;
    }
  }
}
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols iccp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
iccp {
  traceoptions;
  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;
  }
}
```

Related Documentation

- iccp on page 704
- *Configuring ICCP for MC-LAG*

[edit protocols lacp] Hierarchy Level

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>);
  }
fast-hello-issu
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols lldp] Hierarchy Level

```
protocols {
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier seconds;
    interface (all | interface-name) {
      disable;
    }
    lldp-configuration-notification-interval seconds;
    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;
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols oam] Hierarchy Level

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
      event {
        adjacency-loss;
        interface-status-tlv (down | lower-layer-down);
        port-status-tlv blocked;
        rdi;
```

 Copyright © 2015, Juniper Networks, Inc.

```
    }
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    bridge-domain name;
    routing-instance r1 {
    bridge-domain name;
    instance vpls-instance;
    interface (ge | xe) fpc/pic/port.domain;
    level number;
    maintenance-association name{
      mep identifier {
        direction (up | down)
        interface (ge | xe) fpc/pic/port.domain (working | protect );
        auto-discovery;
        lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
          rem-err-xcon | xcon);
        priority number;
      }
    }
    mip-half-function (none | default | explicit);
    name-format (character-string | none | dns | mac+2oct);
    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;
      }
      maintenance-association ma-name {
      mip-half-function (none | default | explicit);
      mep mep-id {
        auto-discovery;
        direction (up | down);
        interface interface-name (working |protect);
        priority number;
        remote-mep mep-id {
          action-profile profile-name;
          sla-iterator-profile profile-name {
            data-tlv-size bytes;
            iteration-count frames;
            priority priority-value;
          }
        }
      }
    }
  }
  performance-monitoring {
    hardware-assisted-timestamping;
    sla-iterator-profiles {
      profile-name {
        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);
            }
        }
    }
    no-aggregate-delegate-processing;
}
link-fault-management {
    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;
        }
    }
    interface interface-name {
        apply-action-profile profile-name;
        event-thresholds {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
        }
        link-discovery (active | passive);
        negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
        }
        pdu-interval interval;
        pdu-threshold threshold-value;
        remote-loopback;
    }
}
fnp {
    interval <100ms | 1s | 10s | 1m | 10m>;
    loss-threshold number
    interface interface name {
            domain-id domain-id
    }
}
}
```

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols ppp] Hierarchy Level

```
monitor-session (interface-name | all);
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable> ;
  flag flag <disable>;
}
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols pppoe] Hierarchy Level

```
pppoe {
  no-send-pads-error;
  no-send-pads-ac-info
  pado-advertise;
  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;
        }
      }
    }
  }
  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;
    }
  }
```

[edit protocols protection-group] Hierarchy Level

```
ethernet-ringring-name {
  east-interface {
    control-channel channel-name {
      vlan number;
    }
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
  west-interface {
    control-channel channel-name {
      vlan number;
    }
  }
}
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols vrrp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  vrrp {
    asymmetric-hold-time;
    delegate-processing;
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
      }
    version-3;
  }
```

 Copyright © 2015, Juniper Networks, Inc.

```
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces Feature Guide for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit system processes] Hierarchy Level

```
dialer-services {
  disable;
}
isdn-signaling {
  disable;
  reject-incoming;
}
```

Related Documentation

- *ISDN Configuration Overview*
- *Disabling ISDN Processes*

Copyright © 2015, Juniper Networks, Inc.

802.3ad

Syntax

```
802.3ad {
  ae interface-number (primary | backup);
  lacp {
    port-priority;
  }
}
```

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

ae ***interface-number***—Aggregated Ethernet logical interface number.

Range: 0 through 15

primary—For link protection configurations, specify the primary link for egress traffic.

backup—For link protection configurations, specify the backup link for egress traffic.

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 42
- Configuring Aggregated Ethernet Link Protection on page 63

Copyright © 2015, Juniper Networks, Inc.

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 statements are explained separately.

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.

Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Configuring MAC Address Filtering on page 316
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 18
- **source-filtering on page 852**

Copyright © 2015, Juniper Networks, Inc.

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 (Intelligent Queuing 2 (IQ2) PICs on M120 and M320 routers; MPCs on MX Series routers) 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 201
- *Configuring the PPPoE Family for an Underlying Interface*
- *Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles*
- PPPoE Overview on page 188

Copyright © 2015, Juniper Networks, Inc.

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 306
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 308
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 309
- [edit chassis] Hierarchy Level on page 570

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 438

 Copyright © 2015, Juniper Networks, Inc.

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 Maintenance Endpoint on page 371

Copyright © 2015, Juniper Networks, Inc.

action-profile (Defining for CFM)

Syntax

```
action-profile profile-name {
  event {
    adjacency-loss;
    interface-status-tlv (down | lower-layer-down);
    port-status-tlv blocked;
    rdi;
  }
  action {
    interface-down;
  }
  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.

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

Related Documentation

- Configuring a Connectivity Fault Management Action Profile on page 376
- default-actions on page 647
- *event (CFM)*
- interface-down on page 719

 Copyright © 2015, Juniper Networks, Inc.

action-profile (Defining for LFM)

Syntax

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

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

Release Information Statement introduced in Junos OS Release 8.5.

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

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

The remaining statements are explained separately.

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

Related Documentation

- Configuring an OAM Action Profile on page 437

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.

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 73
- Example: Configuring Aggregated Ethernet Load Balancing on page 78

 Copyright © 2015, Juniper Networks, Inc.

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;
  (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*]

Copyright © 2015, Juniper Networks, Inc.

Release Information Statement introduced before Junos OS Release 7.4.
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 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, 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 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/24;
     }
  }
}
```

For more information on configuring the same address on multiple interfaces, see *Configuring the Interface Address*.

The remaining statements are explained separately.

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.

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- *Configuring the Protocol Family*
- *Junos OS Administration Library for Routing Devices*
- *family*
- negotiate-address on page 774
- *unnumbered-address (Ethernet)*

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.

Description For MX Series and T Series routers and EX Series switches, 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 178
- *show lldp*
- *Configuring LLDP (CLI Procedure)*
- *Understanding 802.1X and LLDP and LLDP-MED*
- *transmit-delay*
- *Understanding LLDP*

 Copyright © 2015, Juniper Networks, Inc.

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 379

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.

Copyright © 2015, Juniper Networks, Inc.

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 on page 206
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 211

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 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 Gigabit Ethernet Policers on page 311
- *premium (Hierarchical Policer)*
- ieee802.1p on page 705

Copyright © 2015, Juniper Networks, Inc.

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 48

 Copyright © 2015, Juniper Networks, Inc.

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 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 Interfaces Overview on page 3

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 262
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294

 Copyright © 2015, Juniper Networks, Inc.

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 253

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 444

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 442

Copyright © 2015, Juniper Networks, Inc.

arp (Interfaces)

Syntax arp *ip-address* (mac | multicast-mac) *mac-address* publish;

Hierarchy Level [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*]

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 Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.

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, **0011.2233.4455** or **00:11:22:33:44:55**.

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, **0011.2233.4455** or **00:11:22:33:44:55**.

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: 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 Static ARP Table Entries on page 229

Copyright © 2015, Juniper Networks, Inc.

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 Access Routers.

Description For all 10-Gigabit Ethernet interfaces, M120, M320, and T Series routers, 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

- 10-Gigabit Ethernet Notification of Link Down Alarm Overview on page 261

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 622**
- **dot1x on page 655**

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.

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 621
- dot1x on page 655

 Copyright © 2015, Juniper Networks, Inc.

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 Maintenance Endpoint on page 371

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 8.4 for J Series Services Routers.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access 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 and J Series Services Routers, 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 J Series Services Routers with universal Physical Interface Modules (uPIMs) and 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.

 Copyright © 2015, Juniper Networks, Inc.

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 323
- Configuring Gigabit Ethernet Interfaces on J Series Services Routers on page 6
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

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 For J Series Services Routers with 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 202
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2015, Juniper Networks, Inc.

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 preemptive fast reroute (FRR) insertion.

Default By default, FRR insertion is disabled.

Options **backward-frr-enable**—Enable backward FRR insertion.

no-backward-frr-enable—Do not enable backward FRR insertion.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 311
- burst-size-limit (Policer for Gigabit Ethernet Interfaces) on page 635

 Copyright © 2015, Juniper Networks, Inc.

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 on page 365
- **maintenance-domain on page 754**

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.

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.

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 295 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 56 on page 628 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 56: Default Clear Threshold Values

PIC	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

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

 Copyright © 2015, Juniper Networks, Inc.

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 56 on page 628 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 295.

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 295
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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.

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 295 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.

 Copyright © 2015, Juniper Networks, Inc.

Table 57 on page 631 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 57: Default Signal Degrade Threshold Values

PIC	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

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.

Copyright © 2015, Juniper Networks, Inc.

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 57 on page 631.

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 295.

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 295
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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 of the source of the BFD session.

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 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).

Default: automatic

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

- *authentication*
- *detection-time*
- *transmit-interval*
- *Configuring Independent Micro BFD Sessions for LAG*
- Example: Configuring Independent Micro BFD Sessions for LAG on page 95
- Understanding Independent Micro BFD Sessions for LAG on page 92

 Copyright © 2015, Juniper Networks, Inc.

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 311
- bandwidth-limit (Policer for Gigabit Ethernet Interfaces) on page 626

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Copyright © 2015, Juniper Networks, Inc.

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

Default: 100 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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 477
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- **delay on page 648**
- **delay-variation on page 650**

chassis

Syntax chassis { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure router chassis properties.

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

Related Documentation

- *Router Chassis Configuration Statements*

Copyright © 2015, Juniper Networks, Inc.

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 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

- Specifying an Output Priority Map on page 313
- input-priority-map on page 712

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 107

 Copyright © 2015, Juniper Networks, Inc.

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 On J Series Services Routers, 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 202

Copyright © 2015, Juniper Networks, Inc.

connectivity-fault-management

Syntax

```
connectivity-fault-management {
  action-profile profile-name {
    default-actions {
      interface-down;
    }
    event {
      adjacency-loss;
      interface-status-tlv (down | lower-layer-down);
      port-status-tlv blocked;
      rdi;
    }
  }
  performance-monitoring {
    delegate-server-processing;
    hardware-assisted-timestamping;
    sla-iterator-profiles {
      profile-name {
        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 {
    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;
```


Copyright © 2015, Juniper Networks, Inc.

```
            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;
              }
            }
          }
        }
        virtual-switch routing-instance-name {
          bridge-domain name <vlan-ids [ vlan-ids ]>;
        }
      }
      no-aggregate-delegate-processing;
    }
```

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.

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 360

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

- Continuity Check Protocol on page 368

Copyright © 2015, Juniper Networks, Inc.

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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

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.

Required Privilege Level

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

Copyright © 2015, Juniper Networks, Inc.

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 477
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

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*
- *Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 855
- Configuring a Remote MEP with an Iterator Profile on page 486

 Copyright © 2015, Juniper Networks, Inc.

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 Connectivity Fault Management Action Profile on page 376

Copyright © 2015, Juniper Networks, Inc.

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 477
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- calculation-weight on page 637

 Copyright © 2015, Juniper Networks, Inc.

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 on page 206

Copyright © 2015, Juniper Networks, Inc.

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 448

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 477
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*
- calculation-weight on page 637

 Copyright © 2015, Juniper Networks, Inc.

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.

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*
- address on page 609
- negotiate-address on page 774
- *Junos OS Administration Library for Routing Devices*

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.

Description Configure the number of aggregated logical devices available to the router.

Options **number**—Number of aggregated logical devices available to the router.

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.

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 48
- *Configuring Aggregated SONET/SDH Interfaces*

 Copyright © 2015, Juniper Networks, Inc.

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 Maintenance Endpoint on page 371
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 360

Copyright © 2015, Juniper Networks, Inc.

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 (MX Series and T Series routers and EX Series switches) 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*.

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 178

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.

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 (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

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.

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 622
- authentication-profile-name on page 621
- interface (IEEE 802.1x) on page 717

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 Routing Between PE and CE Routers in Layer 3 VPNs*

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 on page 206

 Copyright © 2015, Juniper Networks, Inc.

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 on page 206
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

Copyright © 2015, Juniper Networks, Inc.

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 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 Ring Protection Switching Overview on page 115
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- west-interface on page 918
- ethernet-ring on page 675
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*

Copyright © 2015, Juniper Networks, Inc.

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (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*

Copyright © 2015, Juniper Networks, Inc.

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 Access Routers. Only the **atm-ccc-cell-relay** and **atm-ccc-vc-mux** options are supported on ACX Series routers.

Description Configure a logical link-layer encapsulation type.

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.

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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.

 Copyright © 2015, Juniper Networks, Inc.

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 MPLS LSP Tunnel Cross-Connects Using CCC*
- *Circuit and Translational Cross-Connects Overview*
- Identifying the Access Concentrator on page 201
- *Configuring ATM Interface Encapsulation*
- Configuring VLAN Encapsulation on page 148
- Configuring Extended VLAN Encapsulation on page 150
- *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*
- *CoS on ATM IMA Pseudowire Interfaces Overview*
- *Configuring Policing on an ATM IMA Pseudowire*

encapsulation (Physical Interface)

Syntax 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);

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces rlsq *number*:*number*]

Release Information Statement introduced before Junos OS Release 7.4.
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 Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Default **ppp**—Use serial PPP encapsulation.

Options **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.

atm-pvc—Use ATM PVC encapsulation.

cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.

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-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.

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

 Copyright © 2015, Juniper Networks, Inc.

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.

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. 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), 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.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

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. When you use this encapsulation type, you can configure the **ccc** family only.

 Copyright © 2015, Juniper Networks, Inc.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

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.

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-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.

Required Privilege Level

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

Related Documentation

- *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 Encapsulation on page 148
- Configuring Extended VLAN Encapsulation on page 150
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*
- *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 Q-in-Q Tunneling (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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.

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)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

Copyright © 2015, Juniper Networks, Inc.

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 48
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

ethernet (Protocols OAM)

Syntax

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
    }
    performance-monitoring {
      delegate-server-processing;
      hardware-assisted-timestamping;
      sla-iterator-profiles {
        profile-name {
          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);
        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;
```

 Copyright © 2015, Juniper Networks, Inc.

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

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in Junos OS Release 8.2.

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

The remaining statements are explained separately.

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

Related Documentation
- Enabling IEEE 802.3ah OAM Support on page 431

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 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 Gigabit Ethernet Policers on page 311

 Copyright © 2015, Juniper Networks, Inc.

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.

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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 Access 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

 Copyright © 2015, Juniper Networks, Inc.

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.

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 311
- Configuring MAC Address Filtering on page 316
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 325
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

evcs

Syntax

```
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;
}
```

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

evc-protocol cfm | vpls—Specify connectivity fault management (CFM) or virtual private LAN service (VPLS) as the Ethernet Virtual Connection (EVC) protocol.

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.

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.

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 380
- lmi (Ethernet OAM) on page 742

 Copyright © 2015, Juniper Networks, Inc.

event (LFM)

Syntax

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

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

Release Information Statement introduced in Junos OS Release 8.5.

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

The remaining statements are explained separately.

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

Related Documentation

- Monitoring Protocol Status on page 440

event-thresholds

Syntax

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

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

Release Information Statement introduced in Junos OS Release 8.4.

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

The remaining statements are explained separately.

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 434

Copyright © 2015, Juniper Networks, Inc.

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 107

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

Related Documentation
- Configuring VRRP and VRRP for IPv6 on page 184

Copyright © 2015, Juniper Networks, Inc.

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>;
(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;
```

 Copyright © 2015, Juniper Networks, Inc.

```
            (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 stanzas 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
- **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
- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—Multiprotocol Label Switching (MPLS)
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpls**—(M Series and T Series routers only) Virtual private LAN service

The remaining statements are explained separately.

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

Related Documentation

- *Configuring the Protocol Family*

 Copyright © 2015, Juniper Networks, Inc.

fast-aps-switch

Syntax fast-aps-switch;

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

Release Information Statement introduced in Junos OS Release 12.1.

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

NOTE:

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

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

Related Documentation
- *Reducing APS Switchover Time in Layer 2 Circuits*

Copyright © 2015, Juniper Networks, Inc.

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 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 Interfaces Overview on page 3

 Copyright © 2015, Juniper Networks, Inc.

fec

Syntax fec (efec | gfec | gfec-sdfec | 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.
ufec option introduced in Junos OS Release 13.3 for MX Series routers.

Description Enable Forward Error Correction (FEC) mode.

Default If you do not specify a mode, the default mode on PTX Series routers is **gfec-sdfec**.

Options efec—(For M Series, MX Series, and SRX Series routers only) Enhanced Forward Error Correction (EFEC) is configured to detect and correct bit errors.

gfec—(For M Series, MX Series, and SRX Series routers only) G.709 Forward Error Correction (GFEC) mode is configured to detect and correct bit errors.

gfec-sdfec—(For PTX Series routers only) GFEC and Soft Decision Forward Error Correction (SDFEC) modes are configured to detect and correct bit errors.

none—(For M Series, MX Series, and SRX Series routers only) FEC mode is not configured.

ufec—(MX Series routers only) Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors.

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 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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 switches.

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 on page 139
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 141

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 Access 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.

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 13
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

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 458
- Configuring the Failure Notification Protocol on page 522

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 107

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 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

- Specifying an Output Priority Map on page 313
- input-priority-map on page 712
- **forwarding-class** statement in the *Class of Service Feature Guide for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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.

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 278
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4)(Type 4 1X100GE PIC for STFPC4 FPC) Using SA Multicast Mode*
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 282
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 285
- **sa-multicast (100-Gigabit Ethernet) on page 842**
- **vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 906**
- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 907**

Copyright © 2015, Juniper Networks, Inc.

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.

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 287
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 282
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 285

 Copyright © 2015, Juniper Networks, Inc.

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: 1 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 434
- Configuring Threshold Values for Fault Events in an Action Profile on page 442

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: 1 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 434
- Configuring Threshold Values for Fault Events in an Action Profile on page 442

 Copyright © 2015, Juniper Networks, Inc.

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: 1 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 434
- Configuring Threshold Values for Fault Events in an Action Profile on page 442

Copyright © 2015, Juniper Networks, Inc.

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 257
- *Configuring SONET Options for 10-Gigabit Ethernet Interfaces*

 Copyright © 2015, Juniper Networks, Inc.

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>;
  (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 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 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*]

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 Access 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 20
- **no-gratuitous-arp-request on page 776**

 Copyright © 2015, Juniper Networks, Inc.

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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- Enabling the Hardware-Assisted Timestamping Option on page 498

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 before flushing the maintenance association end point (MEP) database, if no updates occur.

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 on page 368
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

 Copyright © 2015, Juniper Networks, Inc.

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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*

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 (MX Series and T Series routers and EX series switches only) 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 178

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.

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- *Configuring ICCP for MC-LAG*

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 313

Copyright © 2015, Juniper Networks, Inc.

igmp-snooping

Syntax

```
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;
  }
}
```

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.

Description Enable IGMP snooping on the router.

 Copyright © 2015, Juniper Networks, Inc.

Default IGMP snooping is disabled on the router.

Options The statements are explained separately.

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

Related Documentation
- *Understanding IGMP Snooping*
- *IGMP Snooping in MC-LAG Active-Active on MX Series Routers 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 Access 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 20

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

Description Add the configured number of bytes to the length of a packet entering the interface.

Options ***bytes***—Number of bytes added to a packet entering 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

- *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 13

 Copyright © 2015, Juniper Networks, Inc.

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 331

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 331

 Copyright © 2015, Juniper Networks, Inc.

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.

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*
- Applying a Policer on page 314
- input-three-color on page 713
- layer2-policer on page 728
- logical-interface-policer on page 745
- output-policer on page 791
- output-three-color on page 793

Copyright © 2015, Juniper Networks, Inc.

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 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

- Specifying an Input Priority Map on page 313
- output-priority-map on page 792

Copyright © 2015, Juniper Networks, Inc.

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.

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*
- Applying a Policer on page 314
- **input-policer on page 711**
- **layer2-policer on page 728**
- **logical-interface-policer on page 745**
- **output-policer on page 791**
- **output-three-color on page 793**

Copyright © 2015, Juniper Networks, Inc.

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.

Description

For aggregated Ethernet interfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP only, define the rewrite profile to be applied to incoming frames on this logical interface.

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Stacking a VLAN Tag on page 335
- output-vlan-map (Aggregated Ethernet) on page 794

Copyright © 2015, Juniper Networks, Inc.

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 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 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

- Stacking a VLAN Tag on page 335
- **output-vlan-map on page 795**
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

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 on page 365
- **maintenance-domain on page 754**

Copyright © 2015, Juniper Networks, Inc.

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

(MX Series and T Series routers and EX Series switches only) 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.

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*.

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 178

 Copyright © 2015, Juniper Networks, Inc.

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.

Copyright © 2015, Juniper Networks, Inc.

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 622
- dot1x on page 655

interface (OAM Link-Fault Management)

Syntax

```
interface interface-name {
  apply-action-profile profile-name;
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
```

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

Release Information Statement introduced in Junos OS Release 8.2.

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

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

The remaining statements are described separately.

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

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 431

 Copyright © 2015, Juniper Networks, Inc.

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 Connectivity Fault Management Action Profile on page 376

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router or switch.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

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

Related Documentation

- *Physical Interface Configuration Statements Overview*
- Configuring Aggregated Ethernet Link Protection on page 63

Copyright © 2015, Juniper Networks, Inc.

interface-mode

Syntax interface-mode (access | trunk);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit interfaces *interface-name* unit *logical-unit-number* family ethernet-switching],
[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 13.2X50-D10 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

NOTE: This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *port-mode*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure interface-mode and irb for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see "Configuring a Trunk Interface on a Bridge Network" on page 156.

Options **access**—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the **vlan-id** statement.

trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the **vlan-id** or **vlan-id-list** statement.

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 Access Mode on page 154
- Configuring a Logical Interface for Trunk Mode on page 155

Copyright © 2015, Juniper Networks, Inc.

- *Example: Connecting Access Switches to a Distribution Switch*

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 115
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

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]

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 398

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.

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 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 295 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.

 Copyright © 2015, Juniper Networks, Inc.

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.

Range: 1 ms through 1000 ms.

Default: 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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 855
- Configuring a Remote MEP with an Iterator Profile on page 486

 Copyright © 2015, Juniper Networks, Inc.

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 477
- *Configuring an Iterator Profile on a Switch (CLI Procedure)*

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 56
- port-priority on page 811

lacp (Aggregated Ethernet)

Syntax

```
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

[edit interfaces aex aggregated-ether-options]
[edit logical-systems *logical-system-name* 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.
fast-failover option introduced in Junos OS Release 12.2.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Configure the Link Aggregation Control Protocol (LACP) for aggregated Ethernet interfaces only.

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.

Default

If you do not specify LACP as either **active** or **passive**, LACP remains passive.

Options

active—Initiate transmission of LACP packets.

Copyright © 2015, Juniper Networks, Inc.

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.

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 56
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 318

 Copyright © 2015, Juniper Networks, Inc.

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 CFM messages to identify the maintenance association.

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

- Configuring the Maintenance Domain Level on page 364
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Copyright © 2015, Juniper Networks, Inc.

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 379

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 440

 Copyright © 2015, Juniper Networks, Inc.

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 432

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 438

Copyright © 2015, Juniper Networks, Inc.

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 442

 Copyright © 2015, Juniper Networks, Inc.

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);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    negotiation-options {
      allow-remote-loopback;
      no-allow-link-events;
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 8.2.

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

The remaining statements are explained separately.

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

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 431

Copyright © 2015, Juniper Networks, Inc.

link-mode

Syntax link-mode *mode*;

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 Access 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, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode. The router's management Ethernet interface, **fxp0** or **em0**, the built-in Fast Ethernet interfaces on the FIC (M7i router), and the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode 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 J Series ePIM Fast Ethernet interfaces, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."

NOTE:

- On EX4300 switches, the interfaces operate in full-duplex mode only.
- 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.

Copyright © 2015, Juniper Networks, Inc.

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 14
- *Understanding Management Ethernet Interfaces*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

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.
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*]**

For LACP standards link protection, specify **link-protection** at the following hierarchy levels:

- For global LACP link protection, specify at **[edit chassis aggregated-devices ethernet lacp]**
- For a specific aggregated Ethernet interface, specify at **[edit interfaces aeX aggregated-ether-options lacp]**

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 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 Aggregated Ethernet Link Protection on page 63
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level [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.

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 routers 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.

 Copyright © 2015, Juniper Networks, Inc.

- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

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 38
- Configuring Aggregated Ethernet Link Speed on page 51
- Configuring Mixed Aggregated Ethernet Links on page 45
- *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 51
- *Configuring Aggregated SONET/SDH Interfaces*

Copyright © 2015, Juniper Networks, Inc.

lldp

Syntax

```
lldp {
  advertisement-interval seconds;
  disable;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  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>;
  }
}
```

Hierarchy Level

[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

(MX Series and T Series routers and EX Series switches only) Specify LLDP configuration parameters.

Options

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring LLDP on page 178

Copyright © 2015, Juniper Networks, Inc.

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 (MX Series and T Series routers and EX Series switches only) 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 178

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.

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Configuring Ethernet Local Management Interface on page 380
- evcs on page 678

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 73
- Example: Configuring Aggregated Ethernet Load Balancing on page 78

Copyright © 2015, Juniper Networks, Inc.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 107

 Copyright © 2015, Juniper Networks, Inc.

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 Tricolor Marking Policers*
- *action*
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 318
- *action*

Copyright © 2015, Juniper Networks, Inc.

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]

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 Access Routers.

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.

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 19

Copyright © 2015, Juniper Networks, Inc.

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 313

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.

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 on page 368
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

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 Endpoint Lowest Priority Defect on page 373

 Copyright © 2015, Juniper Networks, Inc.

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, **0011.2233.4455** or **00:11:22:33:44:55**.

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 168
- *Example: Configuring an IRB Interface*

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 Access 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 316

 Copyright © 2015, Juniper Networks, Inc.

mac-learn-enable

Syntax (mac-learn-enable | no-mac-learn-enable);

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile]

Release Information Statement introduced before Junos OS Release 7.4.

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 configure whether source and destination MAC addresses are dynamically learned:

- **mac-learn-enable**—Allow the interface to dynamically learn source and destination MAC addresses.
- **no-mac-learn-enable**—Prohibit the interface from dynamically learning source and destination MAC addresses.

MAC address learning is based on source addresses. You can start accounting for traffic after there has been traffic sent from the MAC address. Once the MAC address is learned, the frames and bytes transmitted to or received from the MAC address can be tracked.

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 *MX Series Layer 2 Configuration Guide*.

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 316

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.

Description Enable IP and MAC address validation for static Ethernet and IP demux interfaces. Supported on MX Series routers only.

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 131
- *Configuring MAC Address Validation on Static Demux Interfaces*

 Copyright © 2015, Juniper Networks, Inc.

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.

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 367
- Configuring a Maintenance Endpoint on page 371
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

Copyright © 2015, Juniper Networks, Inc.

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.

Description Configure the name of the maintenance domain in IEEE-compliant format.

Options *domain-name*—Name of the maintenance domain.

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.

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- Creating the Maintenance Domain on page 363
- Configuring a Maintenance Endpoint on page 371
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

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 107

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.

For information about scaling values for PPPoE interfaces, access the *Subscriber Management Scaling Values (XLS)* spreadsheet from the Downloads box on the *Junos OS Subscriber Management* pathway page for the current release.

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 on page 212
- Configuring PPPoE Service Name Tables on page 206
- *PPPoE Maximum Session Limit Overview*
- *Configuring an Interface Set of Subscribers in a Dynamic Profile*
- *Subscriber Interfaces and PPPoE Overview*

Copyright © 2015, Juniper Networks, Inc.

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 12.2 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.

Options *maximum-links-limit*—Maximum links limit for aggregated devices.

Range: 16, 32, 64

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 48
- Configuring an Aggregated Ethernet Interface on page 42
- *network-services*

 Copyright © 2015, Juniper Networks, Inc.

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 622**
- **dot1x on page 655**
- **interface (IEEE 802.1x) on page 717**

Copyright © 2015, Juniper Networks, Inc.

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;
      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 12.3R2 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.
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.

Description

Enable multichassis link aggregation groups (MC-LAGs), which enables one device to form a logical LAG interface with two or more other devices.

Options

chassis-id *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.

NOTE: The **prefer-status-control-active** statement can be configured with the **status-control standby** statement to prevent the LACP MC-LAG system ID from reverting to the default LACP

 Copyright © 2015, Juniper Networks, Inc.

system ID on ICCP failure. Use this statement 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** statement does not go into standby mode.

To make the **prefer-status-control-active** statement work with the **status-control standby** statement when an interchassis-link logical interface is configured on a 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 needs to be configured as active.**

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 on MX Series Routers Overview*
- *Configuring Multichassis Link Aggregation on MX Series Routers*
- *Configuring Multichassis Link Aggregation on EX Series Switches*

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 477

 Copyright © 2015, Juniper Networks, Inc.

mep

Syntax

```
mep mep-id {
  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-id—Specify the numeric identifier of the MEP.

Range: 1 through 8191

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 a Maintenance Endpoint on page 371

minimum-links

Syntax minimum-links *number*;

Hierarchy Level
[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*]

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

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 upOn 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.

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 53
- *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*

 Copyright © 2015, Juniper Networks, Inc.

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 the Maintenance Domain on page 363
- maintenance-domain on page 754

Copyright © 2015, Juniper Networks, Inc.

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 245
- Configuring Mixed-Rate Mode Operation on page 250

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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 l2vpn site *site-name*],
[edit routing-instances *routing-instance-name* protocols vpls]

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 Access 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.

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.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configuring the Media MTU*.

Copyright © 2015, Juniper Networks, Inc.

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)

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 Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring the Media MTU*
- *Configuring the MTU for Layer 2 Interfaces*
- *Setting the Protocol MTU*

Copyright © 2015, Juniper Networks, Inc.

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*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

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 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 on EX Series Switches*
- *Example: Configuring IGMP Snooping*
- *Configuring IGMP Snooping (CLI Procedure)*
- *IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview*
- *host-only-interface*
- *show igmp-snooping membership*

Copyright © 2015, Juniper Networks, Inc.

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.

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 on MX Series Routers*
- Configuring Aggregated Ethernet Link Protection on page 63
- Example: Configuring Aggregated Ethernet Link Protection on page 65
- **peer on page 800**

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 367
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

 Copyright © 2015, Juniper Networks, Inc.

native-vlan-id

Syntax native-vlan-id *number*;

Hierarchy Level [edit interfaces ge-*fpc/pic/port*],
[edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
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 Configure mixed tagging support for untagged packets on a port for the following:

- M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging
- MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging
- T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP
- EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces

When the **native-vlan-id** statement is included with the **flexible-vlan-tagging** statement, untagged packets are accepted on the same mixed VLAN-tagged port.

The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID 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.

When the **native-vlan-id** statement is included with the **interface-mode** statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.

Options *number*—VLAN ID number.

Range: (ACX Series routers and EX Series switches) 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 Mixed Tagging Support for Untagged Packets on page 140
- Configuring a Logical Interface for Access Mode on page 154
- *Configuring the Native VLAN Identifier (CLI Procedure)*
- *Understanding Bridging and VLANs on EX Series Switches*

Copyright © 2015, Juniper Networks, Inc.

- flexible-vlan-tagging on page 689
- *Understanding Q-in-Q Tunneling on EX Series Switches*

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*
- address on page 609
- unnumbered-address (PPP) on page 898
- *Junos OS Administration Library for Routing Devices*

negotiation-options

Syntax

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

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

Release Information Statement introduced in Junos OS Release 8.4.

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

The statements are explained separately.

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

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 427

 Copyright © 2015, Juniper Networks, Inc.

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 435

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 Access 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 699.

Copyright © 2015, Juniper Networks, Inc.

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 Access 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 20
- **gratuitous-arp-reply on page 700**

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 205
- *Configuring Frame Relay Keepalives*

 Copyright © 2015, Juniper Networks, Inc.

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 237
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 246

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 on page 215

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 on page 215

 Copyright © 2015, Juniper Networks, Inc.

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.

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 737
- Configuring Aggregated Ethernet Link Protection on page 63
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*

Copyright © 2015, Juniper Networks, Inc.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

oam

Syntax

```
oam {
  ethernet {
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      performance-monitoring {
        delegate-server-processing;
        hardware-assisted-timestamping;
        sla-iterator-profiles {
          profile-name {
            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);
        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.

 Copyright © 2015, Juniper Networks, Inc.

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.

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 427
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

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 rate to Optical Channel Transport Unit 2 (OTU2).

Options **oc192**—OTU2 for 10Gbps

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Copyright © 2015, Juniper Networks, Inc.

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.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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.

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 default threshold values for optical channel data unit (ODU) signal degradation for the P2-100GE-OTN PIC are:

- ber-threshold-clear—1E-09
- ber-threshold-signal-degrade—1E-06
- interval—10 ms

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

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Copyright © 2015, Juniper Networks, Inc.

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.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

optics-options

Syntax

```
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  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.

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.

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 253
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294

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;
  }
  (prbs | no-prbs);
  preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
  }
  rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    otu4;
    (pass-through | no-pass-through);
  }
  signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
  }
  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.

Description

Specify the Ethernet Optical Transport Network (OTN) interface and options.

Options

The remaining statements are explained separately.

Required Privilege Level

interfaces—To view this statement in the configuration.
interfaces-control—To add this statement to the configuration.

 Copyright © 2015, Juniper Networks, Inc.

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 Set the rate to Optical Channel Transport Unit 4 (OTU4).

Default By default, the rate is OTU4 on PTX Series routers.

Options **otu4**—OTU4 for 100Gbps

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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*
- Applying a Policer on page 314
- **input-policer on page 711**
- **input-three-color on page 713**
- **layer2-policer on page 728**
- **logical-interface-policer on page 745**
- **output-three-color on page 793**

Copyright © 2015, Juniper Networks, Inc.

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 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

- Specifying an Output Priority Map on page 313
- input-priority-map on page 712

Copyright © 2015, Juniper Networks, Inc.

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*
- Applying a Policer on page 314
- input-three-color on page 713
- input-policer on page 711
- layer2-policer on page 728
- logical-interface-policer on page 745
- output-policer on page 791

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.

Description

For aggregated Ethernet interfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP only, define the rewrite profile to be applied to outgoing frames on this logical interface.

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 326
- input-vlan-map (Aggregated Ethernet) on page 714

 Copyright © 2015, Juniper Networks, Inc.

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.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

For Gigabit Ethernet IQ, 10-Port 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only, Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite operation to be applied to outgoing frames on this logical interface.

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

- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 326
- **input-vlan-map on page 715**
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

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 on page 206
- Enabling Advertisement of Named Services in PADO Control Packets on page 214

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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 Passive Monitoring on SONET/SDH Interfaces*
- *multiservice-options*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2015, Juniper Networks, Inc.

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 379

pdu-interval

Syntax pdu-interval *interval*;

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

Release Information Statement introduced in Junos OS Release 8.2.

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

Options *interval*—Periodic OAM PDU sending interval.

Range: 100 through 1000 milliseconds

Default: 1000 milliseconds

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

Related Documentation

- Configuring the OAM PDU Interval on page 433

Copyright © 2015, Juniper Networks, Inc.

pdu-threshold

Syntax pdu-threshold *threshold-value*;

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

Release Information Statement introduced in Junos OS Release 8.2.

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

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

Range: 3 through 10 PDUs

Default: 3 PDUs

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

Related Documentation

- Configuring the OAM PDU Threshold on page 433

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 on MX Series Routers*
- Configuring Aggregated Ethernet Link Protection on page 63
- Example: Configuring Aggregated Ethernet Link Protection on page 65
- **multi-chassis-protection on page 771**

Copyright © 2015, Juniper Networks, Inc.

performance-monitoring

Syntax

```
performance-monitoring {
  delegate-server-processing;
  hardware-assisted-timestamping;
  sla-iterator-profiles {
    profile-name {
      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.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 448
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 488
- Enabling the Hardware-Assisted Timestamping Option on page 498

periodic

Syntax periodic *interval*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options lacp],
[edit interfaces interface-range *name* 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.

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 56
- *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*

 Copyright © 2015, Juniper Networks, Inc.

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 407
- policer (CFM Global) on page 804
- policer (CFM Session) on page 805

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 407
- policer (CFM Session) on page 805

 Copyright © 2015, Juniper Networks, Inc.

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 407
- **policer (CFM Global) on page 804**

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.

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 311

 Copyright © 2015, Juniper Networks, Inc.

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 316

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 336
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

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 337

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 338

 Copyright © 2015, Juniper Networks, Inc.

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.

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: 1 through 255

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 (CLI Procedure)*
- *Configuring Aggregated Ethernet LACP*

Copyright © 2015, Juniper Networks, Inc.

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

(MX Series , T Series, and PTX routers only) 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: 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.

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 740
- Configuring LLDP on page 178
- show lldp neighbors on page 1332

Copyright © 2015, Juniper Networks, Inc.

port-status-tlv

Syntax port-status-tlv blocked;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *tlv-action* event]

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 Connectivity Fault Management Action Profile on page 376
- Configuring Remote MEP Action Profile Support on page 398

pp0 (Dynamic PPPoE)

Syntax

```
pp0 {
  unit logical-unit-number {
    keepalives interval seconds;
    no-keepalives;
    pppoe-options {
      underlying-interface interface-name;
      server;
    }
    ppp-options {
      authentication [ authentication-protocols ];
      chap {
        challenge-length minimum minimum-length maximum maximum-length;
      }
      pap;
    }
    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.

 Copyright © 2015, Juniper Networks, Inc.

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 a Basic PPPoE Dynamic Profile*
- *Configuring a PPPoE Dynamic Profile with Additional Options*
- *Configuring Dynamic Authentication for PPP Subscribers*
- For information about creating static PPPoE interfaces, see Configuring PPPoE on page 199

pppoe-options

Syntax

```
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
```

Hierarchy Level
[edit interfaces pp0 unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]

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.

Description
For J Series Services Routers, M120 Multiservice Edge Routers, M320 Multiservice Edge Service Routers, and MX Series Universal Edge Routers with PPP over Ethernet interfaces, configure PPP over Ethernet-specific interface properties.

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 a PPPoE Interface on page 200

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.

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 199 (for static interfaces)
- *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*
- Assigning a Service Name Table to a PPPoE Underlying Interface on page 214

Copyright © 2015, Juniper Networks, Inc.

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 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

- Specifying an Output Priority Map on page 313
- input-priority-map on page 712

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 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 Gigabit Ethernet Policers on page 311
- aggregate (Gigabit Ethernet CoS Policer) on page 615
- ieee802.1p on page 705

Copyright © 2015, Juniper Networks, Inc.

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 855
- Configuring a Remote MEP with an Iterator Profile on page 486

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 Maintenance Endpoint on page 371

 Copyright © 2015, Juniper Networks, Inc.

protection-group

Syntax

```
protection-group {
  ethernet-ringring-name {
    control-channel channel-name {
      vlan number;
      interface name interface-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;
        interface name 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 {
```

Copyright © 2015, Juniper Networks, Inc.

```
                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 115
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 437

Copyright © 2015, Juniper Networks, Inc.

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 (MX Series and T Series routers and EX Series switches only) 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 178

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 (MX Series and T Series routers and EX Series switches only) 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 178

 Copyright © 2015, Juniper Networks, Inc.

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 335
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 336

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

preemptive-fast-reroute

Syntax

```
preemptive-fast-reroute {
  (backward-frr-enable | no-backward-frr-enable);
  (signal-degrade-monitor-enable | no-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 (FRR) options.

Default By default, FRR insertion and signal degrade monitoring are disabled.

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

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 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

- Specifying an Output Priority Map on page 313
- input-priority-map on page 712

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 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 Gigabit Ethernet Policers on page 311
- aggregate (Gigabit Ethernet CoS Policer) on page 615
- ieee802.1p on page 705

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 Maintenance Endpoint on page 371

 Copyright © 2015, Juniper Networks, Inc.

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 Ethernet TCC on page 233
- Example: Configuring an Ethernet TCC or Extended VLAN TCC on page 233
- **remote on page 834**

- *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 227
- *Configuring Proxy ARP (CLI Procedure)*
- *Configuring Proxy ARP (CLI Procedure)*
- *Example: Configuring Proxy ARP on an EX Series Switch*
- Configuring Gratuitous ARP on page 20

 Copyright © 2015, Juniper Networks, Inc.

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 335
- *Configuring Q-in-Q Tunneling (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 336

 Copyright © 2015, Juniper Networks, Inc.

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 622**
- **dot1x on page 655**
- **interface (IEEE 802.1x) on page 717**

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.
oc192 option introduced in Junos OS Release 13.3 for MX Series routers.

Description Specify rate options.

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

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

 Copyright © 2015, Juniper Networks, Inc.

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 655
- interface (IEEE 802.1x) on page 717
- quiet-period on page 831

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 Passive Monitoring on SONET/SDH Interfaces*

Copyright © 2015, Juniper Networks, Inc.

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 Passive Monitoring on SONET/SDH Interfaces*

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 Ethernet TCC on page 233
- Example: Configuring an Ethernet TCC or Extended VLAN TCC on page 233
- proxy on page 827
- *Junos OS VPNs Library for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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 443

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 Maintenance Endpoint on page 371

Copyright © 2015, Juniper Networks, Inc.

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;
  }
}
```

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.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring a Maintenance Endpoint on page 371

 Copyright © 2015, Juniper Networks, Inc.

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: 5 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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 655
- interface (IEEE 802.1x) on page 717
- quiet-period on page 831

 Copyright © 2015, Juniper Networks, Inc.

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.

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 (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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 115
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*

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*

 Copyright © 2015, Juniper Networks, Inc.

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 on page 206
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

Copyright © 2015, Juniper Networks, Inc.

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 282
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 285
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4)(Type 4 1X100GE PIC for STFPC4 FPC) 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 283
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 287
- **forwarding-mode (100-Gigabit Ethernet) on page 693**
- **sa-multicast (PTX Series Packet Transport Routers) on page 843**

 Copyright © 2015, Juniper Networks, Inc.

- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 907**
- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 278

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 283
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 on page 287

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 438

Copyright © 2015, Juniper Networks, Inc.

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 202

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 622**
- **dot1x on page 655**
- **interface (IEEE 802.1x) on page 717**

 Copyright © 2015, Juniper Networks, Inc.

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.

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 on page 206
- Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag on page 210
- Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag on page 208
- Configuring the Action Taken for the Any Service on page 209

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 For J Series Services Routers with 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 202
- *Junos OS Interfaces and Routing Configuration Guide*

 Copyright © 2015, Juniper Networks, Inc.

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 on page 206
- Assigning a Service Name Table to a PPPoE Underlying Interface on page 214
- *Configuring the PPPoE Family for an Underlying Interface*

Copyright © 2015, Juniper Networks, Inc.

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.

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 on page 206
- Creating a Service Name Table on page 207

 Copyright © 2015, Juniper Networks, Inc.

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 367
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 418

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 signal-degrade thresholds.

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

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 preemptive fast reroute (FRR) signal degrade monitoring.

Default By default, FRR signal degrade monitoring is disabled.

Options **no-signal-degrade-monitor-enable**—Do not enable signal degrade monitoring.

signal-degrade-monitor-enable—Enable 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

- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Copyright © 2015, Juniper Networks, Inc.

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. This statement is not supported on the J Series Services Routers.

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 183, 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

- Enabling Ethernet MAC Address Filtering on page 15
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 18
- **source-filtering on page 852**

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

- Enabling Ethernet MAC Address Filtering on page 15
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers on page 18
- **accept-source-mac on page 601**
- **source-address-filter on page 851**

 Copyright © 2015, Juniper Networks, Inc.

speed (Ethernet)

Syntax speed (10m | 100m | 1g | auto | auto-10m-100m);

Hierarchy Level [edit interfaces *interface-name*],
[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 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

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), the built-in Ethernet interfaces on J Series Services Routers, Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode, 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.

Options You can specify the speed as either **10m** (10 Mbps), **100m** (100 Mbps), or on J Series routers with uPIMs installed 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:

- **10m**—10 Mbps
- **100m**—100 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.

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

Related Documentation

- *Configuring the Interface Speed*
- Configuring the Interface Speed on Ethernet Interfaces on page 12
- *Configuring Gigabit Ethernet Autonegotiation*
- Configuring Gigabit Ethernet Interfaces on J Series Services Routers on page 6
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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*
- no-auto-mdix on page 775

 Copyright © 2015, Juniper Networks, Inc.

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 486.

The remaining statements are explained separately.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 477
- Configuring a Remote MEP with an Iterator Profile on page 486
- Verifying the Configuration of an Iterator Profile on page 478
- Managing Iterator Statistics on page 481
- **sla-iterator-profiles on page 856**

sla-iterator-profiles

Syntax

```
sla-iterator-profiles {
  profile-name {
    calculation-weight {
      delay delay-weight;
      delay-variation delay-variation-weight;
    }
    cycle-time milliseconds;
    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 477.

The remaining statements are explained separately.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring an Iterator Profile on page 477
- Configuring a Remote MEP with an Iterator Profile on page 486
- Verifying the Configuration of an Iterator Profile on page 478
- Managing Iterator Statistics on page 481

 Copyright © 2015, Juniper Networks, Inc.

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 Access 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 325
- **vlan-tags (Stacked VLAN Tags) on page 912**

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 on page 206
- Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client on page 213

 Copyright © 2015, Juniper Networks, Inc.

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 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 622
- dot1x on page 655
- interface (IEEE 802.1x) on page 717

Copyright © 2015, Juniper Networks, Inc.

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 622**
- **dot1x on page 655**
- **interface (IEEE 802.1x) on page 717**

 Copyright © 2015, Juniper Networks, Inc.

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 338
- *Configuring Q-in-Q Tunneling (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.

Copyright © 2015, Juniper Networks, Inc.

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 343

 Copyright © 2015, Juniper Networks, Inc.

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 344

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 On a J Series Services Router with multiport Gigabit Ethernet uPIMs installed and operating in access switching mode, only one physical interface is configured for the entire multiport Gigabit Ethernet uPIM. 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.

Related Documentation

- Configuring Gigabit Ethernet Interfaces on J Series Services Routers on page 6

Copyright © 2015, Juniper Networks, Inc.

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 On a J Series Services Router with Ethernet uPIMs installed and operating in access switching mode, 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.

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 on J Series Services Routers on page 6

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: 1 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 434
- Configuring Threshold Values for Fault Events in an Action Profile on page 442

syslog (OAM Action)

Syntax syslog;

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

Release Information Statement introduced in Junos OS Release 8.5.

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

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

Related Documentation

- Specifying the Actions to Be Taken for Link-Fault Management Events on page 438

 Copyright © 2015, Juniper Networks, Inc.

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

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 56

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.

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.

Related Documentation

- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

 Copyright © 2015, Juniper Networks, Inc.

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 Access Routers.
Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

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 329
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

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 331

 Copyright © 2015, Juniper Networks, Inc.

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 on page 206

traceoptions (Protocols LLDP)

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.

 Copyright © 2015, Juniper Networks, Inc.

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 182

Copyright © 2015, Juniper Networks, Inc.

traceoptions (Ethernet Ring Protection)

Syntax

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

Hierarchy Level [edit protocols protection-group]

Release Information 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 trace options for the protection group.

Default Trace options are not set by default. On some EX Series switches, logging of basic ERPS state transitions is set by default. You can configure trace options on those switches to obtain more details than are provided by the default log. See "Understanding Ethernet Ring Protection Switching Functionality" on page 116 for additional information about default logging of the basic state transitions.

Options **file *filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. You can include the following file options:

- **no-stamp**—(Optional) Do not timestamp trace file.
- **no-world-readable**—(Optional) Do not allow any user to read the log file.
- **replace**—(Optional) Replace the trace file rather than appending to it.
- **size**—(Optional) Maximum trace file size (10240..4294967295).
- **world-readable**—(Optional) Allow any user to read the log file.

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 SBC process operations.
- **configuration**—Trace configuration events.
- **debug**—Trace device monitor events.
- **events**—Trace events to the protocol state machine
- **normal**—Trace normal messages.
- **pdu**—Trace RAPS PDU reception and transmission.
- **periodic-packet-management**—Trace periodic packet management state and events.
- **state-machine**—Trace RAPS state machine.
- **timers**—Trace protocol timers.

Copyright © 2015, Juniper Networks, Inc.

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

Related Documentation

- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*

Copyright © 2015, Juniper Networks, Inc.

traceoptions (Interface Process)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

Hierarchy Level [edit interfaces]

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

Description Define tracing operations for the interface process (dcd).

Default If you do not include this statement, no interface-specific tracing operations are performed.

Options **disable**—(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**.

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 **dcd**.

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**
- **change-events**—Log changes that produce configuration events
- **config-states**—Log the configuration state machine changes
- **kernel**—Log configuration IPC messages to kernel
- **kernel-detail**—Log details of configuration messages to kernel

no-world-readable—(Optional) Disallow any user to read the log file.

 Copyright © 2015, Juniper Networks, Inc.

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.

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.

Related Documentation

- *Tracing Operations of the Interface Process*

Copyright © 2015, Juniper Networks, Inc.

traceoptions (LACP)

Syntax

```
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
  }
fast-hello-issu
```

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced in Junos OS Release 7.6.

Description Define tracing operations for the LACP protocol.

Default If you do not include this statement, no LACP protocol tracing operations are performed.

Options

disable—(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**.

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.

 Copyright © 2015, Juniper Networks, Inc.

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 61

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.

Copyright © 2015, Juniper Networks, Inc.

- **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.

Copyright © 2015, Juniper Networks, Inc.

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 on page 206
- Tracing PPPoE Operations on page 221

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 and EX Series switches 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 178

Copyright © 2015, Juniper Networks, Inc.

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 622**
- **dot1x on page 655**
- **interface (IEEE 802.1x) on page 717**

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 (OC) Loss of Frame defect trigger.
- **oc-lom**—OC Loss of Multiframe defect trigger.
- **oc-los**—OC Loss of Signal defect trigger.
- **oc-wavelength-lock**—OC 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 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-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.
- **odu-uas-th**—ODU Unavailable Seconds Threshold defect trigger.
- **opu-ptm**—Optical Channel Payload (OPU) Payload Type Mismatch defect trigger.

 Copyright © 2015, Juniper Networks, Inc.

- **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:

 Copyright © 2015, Juniper Networks, Inc.

- **oc-lof**—Optical Channel (OC) Loss of Frame defect trigger.
- **oc-lom**—OC Loss of Multiframe defect trigger.
- **oc-los**—OC Loss of Signal defect trigger.
- **oc-tsf**—OC TOE security functionality (TSF) defect trigger.
- **oc-wavelength-lock**—OC 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**—(MX Series routers) 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.

Copyright © 2015, Juniper Networks, Inc.

- **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).

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 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

Copyright © 2015, Juniper Networks, Inc.

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 294
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294
- Configuring 100-Gigabit Ethernet OTN Optics on page 299

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 253
- **optics-options on page 787**
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294

Copyright © 2015, Juniper Networks, Inc.

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 For J Series Services Routers, M120 and M320 Multiservice Edge routers, and MX Series Universal Edge Routers with PPP over Ethernet interfaces, 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 Demux Underlying Interface*
- *Configuring a VLAN Demux Underlying Interface*
- *Specifying the Demux Underlying Interface*
- Configuring the PPPoE Underlying Interface on page 201
- *Junos OS Interfaces and Routing Configuration Guide*

Copyright © 2015, Juniper Networks, Inc.

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);
  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;
  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;
```


Copyright © 2015, Juniper Networks, Inc.

```
  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;
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;
   vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
   chap {
      access-profile name;
      default-chap-secret name;
      local-name name;
      passive;
   }
   compression {
      acfc;
      pfc;
   }
   dynamic-profile profile-name;
   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;
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 {
```


Copyright © 2015, Juniper Networks, Inc.

```
          (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.

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.

Copyright © 2015, Juniper Networks, Inc.

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.

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*
- address on page 609
- negotiate-address on page 774
- *Junos OS Administration Library for Routing Devices*

 Copyright © 2015, Juniper Networks, Inc.

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

Related Documentation

- *Junos OS Support for VRRPv3*
- *VRRP for IPv4 Configuration Hierarchy*
- *VRRP for IPv6 Configuration Hierarchy*

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 on page 366

Copyright © 2015, Juniper Networks, Inc.

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
- Configuring Mixed Tagging on page 139

 Copyright © 2015, Juniper Networks, Inc.

vlan-id (VLAN ID to Rewrite)

Syntax 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 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, 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.

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 338
- Binding VLAN IDs to Logical Interfaces on page 142

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 {
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 142
- encapsulation (Logical Interface) on page 660
- encapsulation (Physical Interface) on page 664
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*
- flexible-vlan-tagging on page 689
- vlan-tagging on page 908
- vlan-tags (Dual-Tagged Logical Interface) on page 910

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.

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 155
- Configuring the VLAN ID List for a Trunk Interface on page 155

Copyright © 2015, Juniper Networks, Inc.

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 on page 143

 Copyright © 2015, Juniper Networks, Inc.

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.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

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 343

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 278
- **forwarding-mode (100-Gigabit Ethernet) on page 693**
- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 907**

 Copyright © 2015, Juniper Networks, Inc.

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 278
- **forwarding-mode (100-Gigabit Ethernet) on page 693**
- **sa-multicast (100-Gigabit Ethernet) on page 842**
- **vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 906**

Copyright © 2015, Juniper Networks, Inc.

vlan-tagging

Syntax vlan-tagging;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* 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 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.

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: 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.

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

Related Documentation

- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *Example: Configuring BGP Autodiscovery for LDP VPLS*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- Configuring Tagged Aggregated Ethernet Interfaces on page 54
- *Configuring Interfaces for VPLS Routing*
- *Enabling VLAN Tagging*
- 802.1Q VLANs Overview on page 136
- *vlan-id*

Copyright © 2015, Juniper Networks, Inc.

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 MX Series routers with 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.

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 331
- **stacked-vlan-tagging on page 857**

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.

 Copyright © 2015, Juniper Networks, Inc.

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 142
- **encapsulation (Logical Interface) on page 660**
- **encapsulation (Physical Interface) on page 664**
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*.
- **flexible-vlan-tagging on page 689**
- **vlan-id-list (Ethernet VLAN Circuit) on page 902**
- **vlan-tagging on page 908**

Copyright © 2015, Juniper Networks, Inc.

vlan-tags (Stacked VLAN Tags)

Syntax vlan-tags inner *tpid.vlan-id* 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 For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.

NOTE: The **inner-range** *vid1–vid2* option is supported on MX Series with IQE PICs only.

Options **inner** ***tpid.vlan-id***—A TPID and a valid VLAN identifier.

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-range ***vid1–vid2***—For MX Series routers with Enhanced IQ (IQE) PICs only; specify a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

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. 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 (1–4094), 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;
}
```

 Copyright © 2015, Juniper Networks, Inc.

```
[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 331
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 141
- **stacked-vlan-tagging on page 857**

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 262
- **optics-options on page 787**
- 100-Gigabit Ethernet OTN Options Configuration Overview on page 294

Copyright © 2015, Juniper Networks, Inc.

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

 Copyright © 2015, Juniper Networks, Inc.

- **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

Copyright © 2015, Juniper Networks, Inc.

- **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

Copyright © 2015, Juniper Networks, Inc.

- **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 253
- Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength on page 253
- **show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet) on page 1171**

Copyright © 2015, Juniper Networks, Inc.

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;
}
```

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 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 Ring Protection Switching Overview on page 115
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- **east-interface on page 658**
- **ethernet-ring on page 675**
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2015, Juniper Networks, Inc.

CHAPTER 33

Operational Commands

- show interfaces (10-Gigabit Ethernet)
- show interfaces (M Series, MX Series and T Series Routers, and PTX Series Packet Transport Routers 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 (Gigabit Ethernet)
- show interfaces mc-ae
- 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

Copyright © 2015, Juniper Networks, Inc.

- 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 data-channel
- 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 vrrp
- traceroute ethernet

Copyright © 2015, Juniper Networks, Inc.

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 922

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
```

 Copyright © 2015, Juniper Networks, Inc.

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 1300

List of Sample Output clear interfaces interval on page 923

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
```

Copyright © 2015, Juniper Networks, Inc.

```
user@host> clear interfaces interval t3-0/3/0:4
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

clear lldp neighbor

Syntax

clear lldp neighbor
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

On MX Series and T Series routers, 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*.

Options

interface ***interface-name***—(Optional) Clear the LLDP neighbors on the specified interface.

Required Privilege Level

clear

Related Documentation

• clear lldp statistics on page 927

List of Sample Output

clear lldp statistics on page 926

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 statistics

```
user@host> clear lldp statistics
user@host> clear lldp statistics interface ge-0/2/0
```

 Copyright © 2015, Juniper Networks, Inc.

clear lldp statistics

Syntax

clear lldpp neighbor
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

On MX Series and T Series routers, 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*.

Options

interface *interface-name*—(Optional) Clear LLDP statistics on the specified interface.

Required Privilege Level

clear

Related Documentation

- clear lldp neighbor on page 926

List of Sample Output

clear lldp neighbor on page 927

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 neighbor

`user@host>` **clear lldp neighbors**

`user@host>` **clear lldp neighbors interface ge-0/2/2**

Copyright © 2015, Juniper Networks, Inc.

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 514
- *Ethernet Interfaces Feature Guide for Routing Devices*

List of Sample Output

clear oam ethernet connectivity-fault-management continuity-measurement on page 928

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.
```


Copyright © 2015, Juniper Networks, Inc.

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.
```

Copyright © 2015, Juniper Networks, Inc.

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 514
- Displaying ETH-LM Statistics on page 513
- Managing ETH-LM Statistics on page 513

List of Sample Output

clear oam ethernet connectivity-fault-management loss-statistics on page 930

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
```

 Copyright © 2015, Juniper Networks, Inc.

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 1390

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
```

Copyright © 2015, Juniper Networks, Inc.

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*
- show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 1398

List of Sample Output clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 932

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.
```


Copyright © 2015, Juniper Networks, Inc.

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 933

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
```

Copyright © 2015, Juniper Networks, Inc.

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 934

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
```

 Copyright © 2015, Juniper Networks, Inc.

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 935
clear protection-group ethernet-ring statistics on page 935

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***

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*>
<wait *interval-seconds*>
<priority *802.1p value*>
<size>
<no-session-id-tlv>
<xml>

Release Information Command introduced in Junos OS Release 9.5.

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 run-time output of the **monitor ethernet delay-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet frame delay measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay—The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

Options **maintenance-domain** ***md-name***—Name of an existing CFM maintenance domain.

maintenance-association ***ma-name***—Name of an existing CFM maintenance association.

 Copyright © 2015, Juniper Networks, Inc.

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.

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 939
monitor ethernet delay-measurement two-way on page 939
monitor ethernet delay-measurement two-way (Invalid DMR Frames Received) on page 940

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 58 on page 938 lists the run-time output fields for the **monitor ethernet delay-measurement one-way** command.
- Table 59 on page 939 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 58: 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.
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.

Copyright © 2015, Juniper Networks, Inc.

Table 58: monitor ethernet delay-measurement one-way Output Fields *(continued)*

Output Field Name	Output Field Description
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 59: 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 *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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

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*>
<wait *interval-seconds*>
<priority *802.1p value*>
<no-session-id-tlv>
<xml>

Release Information

Command introduced in Junos OS Release 11.1.

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.

Copyright © 2015, Juniper Networks, Inc.

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.

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 454
- *Junos OS Network Interfaces Library for Routing Devices*
- CLI Explorer

 Copyright © 2015, Juniper Networks, Inc.

List of Sample Output monitor ethernet loss-measurement on page 944

Output Fields Table 60 on page 943 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 60: 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.
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.

Copyright © 2015, Juniper Networks, Inc.

Table 60: monitor ethernet loss-measurement output fields *(continued)*

Output Field Name	Output Field Description
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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%
```

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 946
request interface mc-ae switchover mcae-id on page 946

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 948

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**

 Copyright © 2015, Juniper Networks, Inc.

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 949

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
```

Copyright © 2015, Juniper Networks, Inc.

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>
```

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, MX2010, and MX2020

```
show chassis hardware
<clei-models>
<detail | extensive>
```

 Copyright © 2015, Juniper Networks, Inc.

3D Universal Edge Routers)

<models>

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 Access 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 for the TX Matrix Plus router in Junos OS Release 9.6.
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 Access 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 MX104 3D Universal Edge Routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

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, 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.

 Copyright © 2015, Juniper Networks, Inc.

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.

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:

Table 61: 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.

Required Privilege Level

view

Related Documentation

- *show chassis power*

List of Sample Output

show chassis hardware (EX8216 Switch) on page 959
show chassis hardware clei-models (EX8216 Switch) on page 960
show chassis hardware clei-models (T1600 Router) on page 961
show chassis hardware detail (EX4200 Switch) on page 962
show chassis hardware (EX4300 Switch) on page 962
show chassis hardware models (EX4500 Switch) on page 962
show chassis hardware detail (EX9200 Switch) on page 962
show chassis hardware (J6350 Router) on page 963

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Output Fields Table 62 on page 957 lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

 Copyright © 2015, Juniper Networks, Inc.

Table 62: 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. See *EX Series Switches Hardware and CLI Terminology Mapping*. • (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 and MX2020 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.	All levels
Version	Revision level of the chassis component.	All levels
Part number	Part number of the chassis component.	All levels
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

Copyright © 2015, Juniper Networks, Inc.

Table 62: show chassis hardware Output Fields *(continued)*

Field Name	Field Description	Level of Output
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. On the J Series routers, the FPC type corresponds to the Physical Interface Module (PIM). 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) • **4x GE Base PIC**—Four built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM) • **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) • **1x ADSL Annex B**—ADSL 2/2+ Annex B PIM (one port, for ISDN)	All levels

 Copyright © 2015, Juniper Networks, Inc.

Table 62: show chassis hardware Output Fields *(continued)*

Field Name	Field Description	Level of Output
	- **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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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**

 Copyright © 2015, Juniper Networks, Inc.

```
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-1OC192-SON-SR2
SIB 0            REV 07   710-013074                     SIB-I-T1600-S
SIB 1            REV 07   710-013074                     SIB-I-T1600-S
```

Copyright © 2015, Juniper Networks, Inc.

```
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 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 (J6350 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1090E07ADB      JSR6350
Midplane         REV 03   710-014593   NP1265
System IO        REV 01   710-016210   NN9950            JX350 System IO
Crypto Module                                            Crypto Acceleration
Routing Engine   REV 08   710-015273   NM6509            RE-J6350-3400
   ad0      248 MB   256MB CKS          00102006C24A00000039 Compact
Flash
FPC 0                                                    FPC
   PIC 0                                                  4x GE Base PIC
FPC 1            REV 06   750-010355   AI07030023        FPC
   PIC 0                                                  2x T1
FPC 3            REV 06   750-011148   AJ06520151        FPC
   PIC 0                                                  2x E1
FPC 6            REV 06   750-013492   NC4170            FPC
   PIC 0                                                  4x FE
Power Supply 0
```

show chassis hardware (J6300 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN000164AB        J6300
Midplane         REV 02.04 710-010001  CORE99570
System IO        REV 02.00 710-010003  CORE100848        System IO board
Routing Engine   RevX2.6   750-010006  IWGS40735390      RE-J.3
FPC 0                                                    FPC
  PIC 0                                                  2x FE
FPC 1            RevX2.0   750-011380  N3960005          FPC
  PIC 0                                                  1xADSL pic Annex A
FPC 2            RevX2.0   750-011380  N3960002          FPC
  PIC 0                                                  1xADSL pic Annex B
FPC 3            REV 03    750-010354  N0780028          FPC
  PIC 0                                                  1x T3
```

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 (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
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
 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 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 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   AARCHOO           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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
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   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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
    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
```

Copyright © 2015, Juniper Networks, Inc.

```
      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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

```
      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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
      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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
 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
```

Copyright © 2015, Juniper Networks, Inc.

```
 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
 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
```

Copyright © 2015, Juniper Networks, Inc.

```
 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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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 (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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
        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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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:
```


Copyright © 2015, Juniper Networks, Inc.

```
  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           0000000000000000000000
PSM 5            REV 0B   740-033727   VG00015           0000000000000000000000
PSM 6            REV 0B   740-033727   VH00097           0000000000000000000000
PSM 7            REV 0C   740-033727   VJ00151           0000000000000000000000
PSM 8            REV 0C   740-033727   VJ00149           0000000000000000000000
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
PSM 8            REV 0C   740-033727   0000000000       00000000000000000000000
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

```
    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+
```


Copyright © 2015, Juniper Networks, Inc.

```
    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+
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
        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
```


Copyright © 2015, Juniper Networks, Inc.

```
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)

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
 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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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**

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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 (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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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:     ...........
```

Copyright © 2015, Juniper Networks, Inc.

```
 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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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           ÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿÿ
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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:
```


Copyright © 2015, Juniper Networks, Inc.

```
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 201105100048 8D8B8D8B 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
     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   AC0946SOBD1       SFP-LX10
    Xcvr 1       REV 02   740-013111   A281922           SFP-T
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
    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 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
```


Copyright © 2015, Juniper Networks, Inc.

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 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
...
```


Copyright © 2015, Juniper Networks, Inc.

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 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 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
```

Copyright © 2015, Juniper Networks, Inc.

```
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 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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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 (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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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 (QFX3500 Switch running Enhanced Layer 2 Software)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                P3566             QFX3500
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           QFX Routing Engine
FPC 0            REV 16   750-036931   P3566-C           QFX3500-48S4Q
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           48x 10G-SFP+
    Xcvr 12      REV 01   740-030658   AD1125A0438       SFP+-10G-USR
    Xcvr 13      REV 01   740-030658   AD1125A02GN       SFP+-10G-USR
  PIC 1                   BUILTIN      BUILTIN           4x 40G-QSFP+
  PIC 2
  MGMT BRD       REV 10   750-036946   BBAW0328          QFX3500-MGMT-RJ45-AFI
Power Supply 0   Rev 05   740-032091   WA13035           JPSU-650W-AC-AFI
Power Supply 1
Fan Tray 0                                              QFX3500 Fan Tray, Front
 to Back Airflow
Fan Tray 1                                              QFX3500 Fan Tray, Front
 to Back Airflow
Fan Tray 2                                              QFX3500 Fan Tray, Front
 to Back Airflow
```

show chassis hardware (QFX5100 Switch running Enhanced Layer 2 Software)

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                TB3113280048      QFX5100-24Q-2P
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           QFX Routing Engine
FPC 0            REV 02   650-049942   TB3113280048      QFX5100-24Q-2P
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           24x 40G-QSFP
    Xcvr 8       REV 01   740-032986   QA470143          QSFP+-40G-SR4
    Xcvr 14      REV 01   740-032986   QB500525          QSFP+-40G-SR4
  PIC 1          REV 02   611-049555   RR3113310169      QFX-EM-4Q
    Xcvr 0       REV 01   740-032986   QC440904          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-032986   QB240154          QSFP+-40G-SR4
    Xcvr 2       REV 01   740-035085   018110105         QSFP+-40G-LPBK
  PIC 2          REV 02   611-049555   RR3113310209      QFX-EM-4Q
    Xcvr 0       REV 01   740-032986   QB190270          QSFP+-40G-SR4
    Xcvr 1       REV 01   740-035085   018110063         QSFP+-40G-LPBK
    Xcvr 2       REV 01   740-032986   QB210034          QSFP+-40G-SR4
Power Supply 0   REV 03   740-041741   1GA23110973        JPSU-650W-AC-AFO
Power Supply 1   REV 03   740-041741   1GA23090878        JPSU-650W-AC-AFO
Fan Tray 0                                            QFX5100 Fan Tray 0, Front
 to Back Airflow - AFO
Fan Tray 1                                            QFX5100 Fan Tray 1, Front
 to Back Airflow - AFO
Fan Tray 2                                            QFX5100 Fan Tray 2, Front
 to Back Airflow - AFO
Fan Tray 3                                            QFX5100 Fan Tray 3, Front
 to Back Airflow - AFO
```

```
                        Fan Tray 4                                   QFX5100 Fan Tray 4, Front
                         to Back Airflow - AFO
```

 Copyright © 2015, Juniper Networks, Inc.

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)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*
<all-members>
<local>
<member *member-id*>

Syntax (MX104, MX2010 and MX2020 3D Universal Edge Routers)

show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Syntax (PTX Series Packet Transport Router)

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 Access 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 Access 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.

Copyright © 2015, Juniper Networks, Inc.

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**.
- 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.

 Copyright © 2015, Juniper Networks, Inc.

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 (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, QFX3600, and OCX Series standalone switches , replace ***slot-number*** with 0.

transport—Display PIC information for optical transport network.

Required Privilege Level

view

Related Documentation

- *request chassis pic*
- show chassis hardware on page 950
- *Configuring the PIC Type*

- 100-Gigabit Ethernet Type 4 PIC with CFP Overview on page 271

List of Sample Output

 Copyright © 2015, Juniper Networks, Inc.

show chassis pic FPC slot PIC slot (MX Routers with Media Services Blade [MSB]) on page 1140
show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers) on page 1140

Output Fields Table 63 on page 1129 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 63: 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.**
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

Copyright © 2015, Juniper Networks, Inc.

Table 63: show chassis pic Output Fields *(continued)*

Field Name	Field Description

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.
Multirate Mode	Rate-selectabilty status for the MIC: **Enabled** or **Disabled**.
Channelization	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.

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
```


Copyright © 2015, Juniper Networks, Inc.

```
  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 pic-slot (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 (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 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
```

 Copyright © 2015, Juniper Networks, Inc.

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 (MX480 Routers 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 routers with OTN Interfaces)

```
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 Routers 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 Routers)

```
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 Routers)

```
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 Routers 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
```


Copyright © 2015, Juniper Networks, Inc.

```
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 Routers 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 Routers 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 (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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 pic-slot (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 pic-slot (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 pic-slot (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 pic-slot (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 pic-slot (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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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 pic-slot (ACX2000 Universal Access 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 pic-slot (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 PIC slot (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
```


Copyright © 2015, Juniper Networks, Inc.

show iccp

Syntax

show iccp <brief | detail>
logical-system [*system-name* | all]

Release Information

Command introduced in Junos OS Release 10.0 for the MX Series.
Support for logical systems added in Junos OS Release 14.1.

Description

Display Interchassis Control Protocol (ICCP) information about the multichassis link aggregation group (MC-LAG) peers, including the state of the TCP connection, Bidirectional Forwarding Detection protocol, backup liveness peer status, and MCSNOOPD, LACPD, and ESWD applications.

Options

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

none—Display ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

brief—Display brief ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

detail—Display detailed ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

Required Privilege Level

view

Related Documentation

- *Understanding Multichassis Link Aggregation*

List of Sample Output

show iccp on page 1142

Output Fields

Table 64 on page 1141 lists the output fields for the **show iccp** command. Output fields are listed in the approximate order in which they appear.

Table 64: show iccp

Field Name	Field Description
Redundancy Group Information for peer	Aggregated Ethernet interface name.
TCP Connection	Specifies if the TCP connection between the peers hosting the MC-LAG is up or down.
Status	Displays the state of the redundancy group: up or down
Redundancy Group ID	Displays the redundancy group identifier that is used to associate with multiple chassis that perform similar failover operations.

Copyright © 2015, Juniper Networks, Inc.

Table 64: show iccp *(continued)*

Field Name	Field Description
Status	Displays the state of the redundancy group: up or down
Client Application	Specifies information regarding the state of the MCSNOOPD and ESWD client applications.
Redundancy Group IDs Joined	Denotes the redundancy group unique identifier that is associated for the particular client application or process.

Sample Output

show iccp

```
user@switch> show iccp
Logical system :LS1
    Redundancy Group Information for peer 16.1.1.1
       TCP Connection       : Established
       Liveliness Detection : Up
       Redundancy Group ID          Status
         2                           Up
         12                         Up

Client Application: lacpd
Redundancy Group IDs Joined:  1
Redundancy Group IDs Joined:  2

Client Application: l2ald_iccpd_client
 Redundancy Group IDs Joined:  1
 Redundancy Group IDs Joined:  2
```


Copyright © 2015, Juniper Networks, Inc.

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 1148
show interfaces brief (Adaptive Services) on page 1148
show interfaces detail (Adaptive Services) on page 1148
show interfaces extensive (Adaptive Services) on page 1149

Output Fields Table 65 on page 1143 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 65: 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
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none

Copyright © 2015, Juniper Networks, Inc.

Table 65: 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**
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**

 Copyright © 2015, Juniper Networks, Inc.

Table 65: 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. • **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
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 65: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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

 Copyright © 2015, Juniper Networks, Inc.

Table 65: 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

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.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.

Description

(M Series, T Series, MX Series, and PTX Series routers and EX Series switches) 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 about the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Aggregated Ethernet) on page 1156
show interfaces brief (Aggregated Ethernet) on page 1157
show interfaces detail (Aggregated Ethernet) on page 1157
show interfaces extensive (Aggregated Ethernet) on page 1158
show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 1159

Output Fields

Table 66 on page 1151 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 66: show interfaces (Aggregated Ethernet) 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
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels

Copyright © 2015, Juniper Networks, Inc.

Table 66: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**

Copyright © 2015, Juniper Networks, Inc.

Table 66: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number of and rate at which bytes and packets are 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 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: • **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**
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**

Copyright © 2015, Juniper Networks, Inc.

Table 66: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
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.	**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

Copyright © 2015, Juniper Networks, Inc.

Table 66: show interfaces (Aggregated Ethernet) 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 load 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.
	• **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.	

Copyright © 2015, Juniper Networks, Inc.

Table 66: show interfaces (Aggregated Ethernet) 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).	
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**

Copyright © 2015, Juniper Networks, Inc.

```
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: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf: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: 10.100.1/24, Local: 10.100.1.2, 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  10.100.1.2/24
```

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: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf: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
```

Copyright © 2015, Juniper Networks, Inc.

```
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: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
      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: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf: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
```

 Copyright © 2015, Juniper Networks, Inc.

```
Logical interface ae0.0 (Index 73) (SNMP ifIndex 563) (Generation 176)
    Flags: Up SNMP-Traps 0x4000 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:
      fe-1/0/3.0
        Input :             0          0             0            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  00:24:dc:85:af:f0       127        2     1

      fe-1/0/3.0   Partner        127  00:23:9c:c3:1f:f0       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-1/0/3.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 224, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.40.1.0/30, Local: 10.40.1.1, Broadcast: 10.40.1.3,
Generation: 187
    Protocol multiservice, MTU: Unlimited, Generation: 225, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (Aggregated Ethernet with VLAN Stacking)

```
user@host> show interfaces ae0 extensive
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: 00:12:1e:19:3f:f0, Hardware address: 00:12:1e:19:3f:f0
  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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 (Demux) on page 1168
show interfaces (PPPoE over Aggregated Ethernet) on page 1169
show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links) on page 1169
show interfaces demux0 (ACI Interface Set Configured) on page 1170

Output Fields Table 67 on page 1162 lists the output fields for the **show interfaces** (demux interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 67: Demux show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
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

 Copyright © 2015, Juniper Networks, Inc.

Table 67: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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**

Copyright © 2015, Juniper Networks, Inc.

Table 67: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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

 Copyright © 2015, Juniper Networks, Inc.

Table 67: Demux show interfaces 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. • **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
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

Table 67: Demux show interfaces Output Fields *(continued)*

<table>
<tr><th>Field Name</th><th>Field Description</th><th>Level of Output</th></tr>
<tr><td>protocol-family</td><td>Protocol family configured on the logical interface.</td><td>brief</td></tr>
<tr><td>Traffic statistics</td><td>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.

</td><td>detail extensive</td></tr>
<tr><td>Local statistics</td><td>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.
</td><td>detail extensive</td></tr>
<tr><td>Transit statistics</td><td>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.
</td><td>detail extensive</td></tr>
<tr><td>IPv6 Transit statistics</td><td>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.
</td><td>detail extensive</td></tr>
<tr><td>Input packets</td><td>Number of packets received on the interface.</td><td>none</td></tr>
</table>

Copyright © 2015, Juniper Networks, Inc.

Table 67: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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

Copyright © 2015, Juniper Networks, Inc.

Table 67: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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 (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: 1.1.1/24
    Traffic statistics:
     Input  bytes  :                    0
```


Copyright © 2015, Juniper Networks, Inc.

```
 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: 11.1.1/24, Local: 11.1.1.1, Broadcast: 11.1.1.255,
    Generation: 434
```

show interfaces (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 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: 100.20.200.202
      Addresses, Flags: Primary Is-Default Is-Primary
        Local: 10.4.12.119
    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
```

 Copyright © 2015, Juniper Networks, Inc.

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*

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 will 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.

The show interfaces diagnostics optics command for optical interfaces does not report the decibel (dBm) value of the received signal if the received power is zero milliwatts (0.0000 mW).

Required Privilege Level view

Related Documentation

- *Supported Network Interface Standards by Transceiver*
- *Supported Network Interface Standards by Transceiver for PTX Series Packet Transport Routers*

List of Sample Output
show interfaces diagnostics optics (DWDM and DWDM OTN) on page 1186
show interfaces diagnostics optics (MPC6E with OTN MIC) on page 1187

show interfaces diagnostics optics (Bidirectional SFP) on page 1188
show interfaces diagnostics optics (SFP) on page 1188
show interfaces diagnostics optics (SFP) on page 1189
show interfaces diagnostics optics (XFP and CFP Optics) on page 1190
show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP) on page 1191
show interfaces diagnostics optics for 40-Gigabit Ethernet on page 1192
show interfaces diagnostics optics for vcp on page 1194

Output Fields Table 68 on page 1172 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 68: 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.
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**.

 Copyright © 2015, Juniper Networks, Inc.

Table 68: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs *(continued)*

Field Name	Field Description
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 69 on page 1173 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 69: 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).
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.
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**.

Copyright © 2015, Juniper Networks, Inc.

Table 69: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
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**.
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: **70.000 mA.**

Copyright © 2015, Juniper Networks, Inc.

Table 69: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
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.**
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.**

Copyright © 2015, Juniper Networks, Inc.

Table 69: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
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 70 on page 1176 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 70: 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.

Copyright © 2015, Juniper Networks, Inc.

Table 70: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(continued)*

Field Name	Field Description
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.
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.

Copyright © 2015, Juniper Networks, Inc.

Table 70: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(continued)*

Field Name	Field Description
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.
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 71 on page 1178 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 71: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.

Copyright © 2015, Juniper Networks, Inc.

Table 71: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(continued)*

Field Name	Field Description
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.
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.

Copyright © 2015, Juniper Networks, Inc.

Table 71: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(continued)*

Field Name	Field Description
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 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.

Copyright © 2015, Juniper Networks, Inc.

Table 71: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(continued)*

Field Name	Field Description
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 72 on page 1181 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 72: 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**.

Copyright © 2015, Juniper Networks, Inc.

Table 72: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(continued)*

Field Name	Field Description
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**.
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**.

Copyright © 2015, Juniper Networks, Inc.

Table 72: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(continued)*

Field Name	Field Description
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.**
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.**

Copyright © 2015, Juniper Networks, Inc.

Table 73 on page 1184 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 73: 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.
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**.

Copyright © 2015, Juniper Networks, Inc.

Table 73: show interfaces diagnostics optics Output for Virtual Chassis Ports *(continued)*

Field Name	Field Description
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**.
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.

Copyright © 2015, Juniper Networks, Inc.

Table 73: show interfaces diagnostics optics Output for Virtual Chassis Ports *(continued)*

Field Name	Field Description
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.
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
        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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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 for vcp

```
ser@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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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 (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 1196
show interfaces far-end-interval coc1-5/2/1:1 on page 1197

Output Fields Table 74 on page 1196 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 74: 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:
...
```


Copyright © 2015, Juniper Networks, Inc.

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***. On the J Series routers, the interface type is **fe-*pim*/0/*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 1211
show interfaces brief (Fast Ethernet) on page 1212
show interfaces detail (Fast Ethernet) on page 1212
show interfaces extensive (Fast Ethernet) on page 1212

Output Fields Table 75 on page 1198 lists the output fields for the **show interfaces Fast Ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 75: 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

 Copyright © 2015, Juniper Networks, Inc.

Table 75: 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

Copyright © 2015, Juniper Networks, Inc.

Table 75: 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 (10-Gigabit Ethernet) command.	**detail extensive**

 Copyright © 2015, Juniper Networks, Inc.

Table 75: 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**

Copyright © 2015, Juniper Networks, Inc.

Table 75: 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.	**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**

Copyright © 2015, Juniper Networks, Inc.

Table 75: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • **Errored blocks**—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode.	**detail extensive**

Copyright © 2015, Juniper Networks, Inc.

Table 75: show interfaces Fast Ethernet Output Fields *(continued)*

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

Copyright © 2015, Juniper Networks, Inc.

Table 75: 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**

Copyright © 2015, Juniper Networks, Inc.

Table 75: 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**

 Copyright © 2015, Juniper Networks, Inc.

Table 75: 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**

Copyright © 2015, Juniper Networks, Inc.

Table 75: 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**

 Copyright © 2015, Juniper Networks, Inc.

Table 75: 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

Copyright © 2015, Juniper Networks, Inc.

Table 75: 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

 Copyright © 2015, Juniper Networks, Inc.

Table 75: 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: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

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  10.10.10.1/24
```

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: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  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: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
        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: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
  Statistics last cleared: Never
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
        Generation: 136
```

Copyright © 2015, Juniper Networks, Inc.

show interfaces (Gigabit Ethernet)

Syntax

show interfaces ge-*fpc/pic/port*
<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 and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.

Options

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.

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

- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

List of Sample Output

show interfaces (Gigabit Ethernet) on page 1229
show interfaces (Gigabit Ethernet on MX Series Routers) on page 1229
show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 1230
show interfaces brief (Gigabit Ethernet) on page 1230
show interfaces detail (Gigabit Ethernet) on page 1231
show interfaces extensive (Gigabit Ethernet IQ2) on page 1232
show interfaces (Gigabit Ethernet Unnumbered Interface) on page 1235
show interfaces (ACI Interface Set Configured) on page 1235

Output Fields

Table 76 on page 1215 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 77 on page 1228.

Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields

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

Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (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**

 Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.	**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 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 76: 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 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.	**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. 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**

 Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**—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**
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**

Copyright © 2015, Juniper Networks, Inc.

Table 76: 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 (10-Gigabit Ethernet) command.
	• **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
	• **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
	• **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
	• **MAC control frames**—Number of MAC control frames.
	• **MAC pause frames**—Number of MAC control frames with **pause** operational code.
	• **Oversized frames**—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**

Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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 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**

Copyright © 2015, Juniper Networks, Inc.

Table 76: 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**

Copyright © 2015, Juniper Networks, Inc.

Table 76: 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**

Copyright © 2015, Juniper Networks, Inc.

Table 76: 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**

 Copyright © 2015, Juniper Networks, Inc.

Table 76: 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

Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
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 Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	**brief detail extensive** 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
Dynamic Profile	(MX Series routers with Trio MPCs only) 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	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	**detail extensive** none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	**detail extensive** none

Copyright © 2015, Juniper Networks, Inc.

Table 76: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) 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**
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	**detail extensive** none

Table 76: show interfaces Gigabit 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 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**

Table 77: 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	

Copyright © 2015, Juniper Networks, Inc.

Table 77: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type *(continued)*

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
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	

Sample Output

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: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  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: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
  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: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
    Egress account overhead: 100
    Ingress account overhead: 90
    Input packets : 0
    Output packets: 0
     Protocol aenet, AE bundle: ae0.0    Link Index: 4
```

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
```

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
```

 Copyright © 2015, Juniper Networks, Inc.

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: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  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
```

Copyright © 2015, Juniper Networks, Inc.

```
      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: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e: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
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

```
    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
```


Copyright © 2015, Juniper Networks, Inc.

```
      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: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
        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 76 on page 1215.

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: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26: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: 22.22.22.22
```

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
```

Copyright © 2015, Juniper Networks, Inc.

```
Input packets : 9
Output packets: 8
Protocol multiservice, MTU: Unlimited
```

 Copyright © 2015, Juniper Networks, Inc.

show interfaces (10-Gigabit Ethernet)

Syntax

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

Release Information

Command introduced in Junos OS Release 8.0.

Description

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

Options

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

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

descriptions—(Optional) Display interface description strings.

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

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

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 1252
show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 1255
show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 1257
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 1259
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 1259
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 1260

Output Fields

See Table 78 on page 1238 for the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields

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

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 78 on page 1238.	**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**

 Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**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**

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**

 Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**

Copyright © 2015, Juniper Networks, Inc.

Table 78: 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 79 on page 1252 • **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**

 Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 78: 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. 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**

 Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

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

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

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Copyright © 2015, Juniper Networks, Inc.

Table 78: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. Table 79 on page 1252 describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In Table 79 on page 1252, 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 79: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

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

Sample Output

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

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


Copyright © 2015, Juniper Networks, Inc.

```
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: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1: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
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
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: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
  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:
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b: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:
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
```

Copyright © 2015, Juniper Networks, Inc.

```
   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
```


Copyright © 2015, Juniper Networks, Inc.

```
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  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: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

Copyright © 2015, Juniper Networks, Inc.

show interfaces (M Series, MX Series and T Series Routers, and PTX Series Packet Transport Routers Management and Internal Ethernet)

List of Syntax

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface) on page 1262
Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface) on page 1262

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

 Copyright © 2015, Juniper Networks, Inc.

List of Sample Output

Output Fields

Table 80 on page 1263 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 Packet Transport Routers. Output fields are listed in the approximate order in which they appear.

Table 80: M Series and T Series Router Management and Internal Ethernet 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 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
Link-level type	Encapsulation type used on the physical interface.	All levels

Table 80: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**

 Copyright © 2015, Juniper Networks, Inc.

Table 80: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields *(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 80: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields *(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:a0:a5:56:01:89, Hardware address: 00:a0:a5:56:01:89
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
        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: 00:80:f9:26:00:c0, Hardware address: 00:80:f9:26:00: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: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
      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:a0:a5:56:01:89, Hardware address: 00:a0:a5:56:01:89
  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: 00:80:f9:26:00:c0, Hardware address: 00:80:f9:26:00: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
```


Copyright © 2015, Juniper Networks, Inc.

```
        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: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
      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:a0:a5:56:01:89, Hardware address: 00:a0:a5:56:01:89
  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: 00:80:f9:26:00:c0, Hardware address: 00:80:f9:26:00: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:
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
      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
```


Copyright © 2015, Juniper Networks, Inc.

```
  Link type      : Full-Duplex
  Current address: 02:00:00:00:00:04, Hardware address: 02:00:00:00:00:04
  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: 02:00:00:22:00:04, Hardware address: 02:00:00:22:00:04
  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: 128/2, Local: 162.0.0.4, Broadcast: 191.255.255.255
    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
```

Copyright © 2015, Juniper Networks, Inc.

```
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: 02:00:00:00:00:04, Hardware address: 02:00:00:00:00:04
  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: 02:00:00:22:00:04, Hardware address: 02:00:00:22:00:04
  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
```


Copyright © 2015, Juniper Networks, Inc.

```
   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: 128/2, Local: 162.0.0.4, 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: 02:00:00:00:00:04, Hardware address: 02:00:00:00:00:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
```

Copyright © 2015, Juniper Networks, Inc.

```
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: 02:00:00:22:00:04, Hardware address: 02:00:00:22:00:04
  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:
```

 Copyright © 2015, Juniper Networks, Inc.

```
 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: 128/2, Local: 162.0.0.4, 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 (J Series Services Routers, 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 1284
show interfaces (PPPoE over Aggregated Ethernet) on page 1284
show interfaces brief (PPPoE) on page 1285
show interfaces detail (PPPoE) on page 1285
show interfaces detail (PPPoE on J Series Services Routers) on page 1286
show interfaces extensive (PPPoE on M120 and M320 Routers) on page 1287

Output Fields Table 81 on page 1278 lists the output fields for the **show interfaces** (PPPoE) command. Output fields are listed in the approximate order in which they appear.

Table 81: 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

Copyright © 2015, Juniper Networks, Inc.

Table 81: 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**

Copyright © 2015, Juniper Networks, Inc.

Table 81: 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		

Copyright © 2015, Juniper Networks, Inc.

Table 81: 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 81: 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

Copyright © 2015, Juniper Networks, Inc.

Table 81: 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

Copyright © 2015, Juniper Networks, Inc.

Table 81: 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: 00:10:94:00:00:01,
      Underlying interface: demux0.100 (Index 88)
    Link:
      ge-1/0/0.32767
      ge-1/0/1.32767
    Input packets : 6
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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: 45.63.24.1
```

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 detail (PPPoE on J Series Services Routers)

```
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
  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 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:
```

 Copyright © 2015, Juniper Networks, Inc.

```
    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: 00:17:cb:48:c8:12,
      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
```

Copyright © 2015, Juniper Networks, Inc.

```
      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: 12.12.12.2, Local: 12.12.12.1, Broadcast: Unspecified,
Generation: 206
```

 Copyright © 2015, Juniper Networks, Inc.

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. Supported in MX Series routers with enhanced queuing DPCs or MPCs.

You can also use the **show interfaces interface-set** command to display information about agent circuit identifier (ACI) interface sets configured on MX Series routers with MPCs/MICs.

Options **interface-set *interface-set-name***—Display information about the specified Gigabit Ethernet, 10-Gigabit Ethernet, or ACI 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*

List of Sample Output

show interfaces interface-set terse on page 1290
show interfaces interface-set detail on page 1290
show interfaces interface-set (ACI Interface Set) on page 1291

Output Fields Table 82 on page 1289 describes the information for the **show interfaces interface-set** command.

Table 82: Ethernet show interfaces interface-set Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Interface set	Name of the interface set or sets.	All levels
Interface set index	Index number of the interface set. For ACI interface sets, the following fields are displayed: • **ACI VLAN**—ACI interface set that the router uses to create dynamic VLAN subscriber interfaces based on the agent circuit identifier value. • **PPPoE**—Dynamic PPPoE subscriber interface that the router creates using the ACI interface set.	**detail** none
Agent Circuit ID	For ACI interface sets, string in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and and the DSL line on the access node.	**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

Table 82: Ethernet show interfaces interface-set Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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 best-effort            211091327             11044380           199995746
    1 expedited-fo                   0                    0                   0
    2 assured-forw                   0                    0                   0
    3 network-cont                   0                    0                   0
  Members:
   xe-11/3/0.0
```

 Copyright © 2015, Juniper Networks, Inc.

show interfaces interface-set (ACI Interface Set)

```
user@host> show interfaces interface-set
Interface set: aci-1001-demux0.1073741826
  Interface set index: 1
    ACI VLAN:
      Agent Circuit ID: aci-ppp-dhcp-dvlan-60
    PPPoE:
      Max Sessions: 3, Max Sessions VSA Ignore: Off
  Members:
   pp0.1073741827
```

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 1293
show interfaces interface-set queue both-ingress-egress (Enhanced DPC) on page 1294
show interfaces interface-set queue egress (Enhanced DPC) on page 1296
show interfaces interface-set queue forwarding-class (Gigabit Ethernet) on page 1297
show interfaces interface-set queue (Enhanced DPC) on page 1298
show interfaces interface-set queue remaining-traffic (Gigabit Ethernet) on page 1298

Output Fields

Table 83 on page 1292 describes the information for the **show interfaces interface-set queue** command.

Table 83: Ethernet show interfaces interface-set queue Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Interface set	Name of the interface set.	All levels
Interface set index	Index number of the interface set.	All levels

Copyright © 2015, Juniper Networks, Inc.

Table 83: Ethernet show interfaces interface-set queue Output Fields *(continued)*

Field Name	Field Description	Level of Output
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:
```


Copyright © 2015, Juniper Networks, Inc.

```
    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
```


Copyright © 2015, Juniper Networks, Inc.

```
     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
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

Copyright © 2015, Juniper Networks, Inc.

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 923

List of Sample Output

show interfaces interval (Channelized OC12) on page 1301
show interfaces interval (E3) on page 1301
show interfaces interval (SONET/SDH) on page 1301

Output Fields

Table 84 on page 1300 lists the output fields for the **show interfaces interval** command. Output fields are listed in the approximate order in which they appear.

Table 84: 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.

Copyright © 2015, Juniper Networks, Inc.

Sample Output

show interfaces interval (Channelized OC12)

```
user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (E3)

```
user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
  17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
  16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

show interfaces interval (SONET/SDH)

```
user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
  20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
  19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
  19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
```

Copyright © 2015, Juniper Networks, Inc.

```
  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:
.....
```


Copyright © 2015, Juniper Networks, Inc.

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 1307
show interfaces irb snmp-index on page 1308

Output Fields Table 85 on page 1303 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 85: 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
Proto	Protocol configured on the interface.	**terse**

Table 85: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**

Copyright © 2015, Juniper Networks, Inc.

Table 85: 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 85: 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	Speed at which the interface is running.	**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

Copyright © 2015, Juniper Networks, Inc.

Table 85: 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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

show interfaces mac-database (Gigabit Ethernet)

Syntax show interfaces mac-database (ge-*fpc/pic/port* | ge-*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.

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 Gigabit Ethernet 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.

mac-address *mac-address*—(Optional) Display detailed MAC address statistics, including policer information.

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 1311
show interfaces mac-database (All MAC Addresses on a Service) on page 1312
show interfaces mac-database mac-address on page 1312

Output Fields Table 86 on page 1309 lists the output fields for the **show interfaces mac-database** command. Output fields are listed in the approximate order in which they appear.

Table 86: 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.

Copyright © 2015, Juniper Networks, Inc.

Table 86: show interfaces mac-database Output Fields *(continued)*

Field Name	Field Description
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.
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.
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 values 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.

Copyright © 2015, Juniper Networks, Inc.

Table 86: show interfaces mac-database Output Fields *(continued)*

Field Name	Field Description

Policer Statistics	(Displayed for **mac-address** option 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
```

Copyright © 2015, Juniper Networks, Inc.

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
```


Copyright © 2015, Juniper Networks, Inc.

show interfaces mc-ae

Syntax

show interfaces mc-ae
<revertive-info>
<id *identifier* unit *number*>

Release Information

Command introduced in Junos OS Release 9.6.
revertive-info statetment introduced in Junos OS Release 13.3

Description

On MX Series routers with multichassis aggregated Ethernet (**ae*X***) interfaces, displaya information about the **ae*X*** interfaces.

Options

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 1314
show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers) on page 1314
show interfaces mc-ae revertive-info on page 1315

Output Fields

Table 87 on page 1313 lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 87: 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**.
Peer Status	Status of the peer link: **active** or **standby**.
Local State	Up or down state of the local device.

Copyright © 2015, Juniper Networks, Inc.

Table 87: show interfaces mc-ae Output Fields *(continued)*

Output Field Name	Field Description
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.

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
 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 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 1316

Output Fields Table 88 on page 1316 describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

Table 88: 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,
```


Copyright © 2015, Juniper Networks, Inc.

```
                        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
```

Copyright © 2015, Juniper Networks, Inc.

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.

Description

(MX Series routers only) Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregate Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or logical systems on a router. Note that the Redundancy Group ID is inherited by the bridging domain 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 1313
- *Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers*

List of Sample Output

show l2-learning redundancy-groups arp-statistics on page 1320
show l2-learning redundancy-groups nd-statistics on page 1320
show l2-learning redundancy-groups remote-macs on page 1320
show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems) on page 1321
show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems) on page 1321
show l2-learning redundancy-groups group-id on page 1321

Copyright © 2015, Juniper Networks, Inc.

show l2-learning redundancy-groups logical-system on page 1321

Output Fields Output fields are listed in the approximate order in which they appear.

Table 89: 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 90: 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.
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).

Copyright © 2015, Juniper Networks, Inc.

Table 91: 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
```


Copyright © 2015, Juniper Networks, Inc.

```
Flags   Status
10          1.1.1.2      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
```

Copyright © 2015, Juniper Networks, Inc.

show lacp interfaces

Syntax

show lacp interfaces
<*interface-name*>

Release Information

Command introduced in Junos OS Release 7.6.

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***

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

List of Sample Output

show lacp interfaces (Aggregated Ethernet) on page 1325
show lacp interfaces (Gigabit Ethernet) on page 1325

Output Fields

Table 64 on page 1141 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 92: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Copyright © 2015, Juniper Networks, Inc.

Table 92: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

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.

Table 92: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

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.

LACP Statistics	LACP statistics are returned when the **extensive** option is used and provides the following information: • **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.

 Copyright © 2015, Juniper Networks, Inc.

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
 Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout    Activity

      ge-1/0/1       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/1     Partner    No   Yes    No   No   No   Yes     Fast     Passive

      ge-1/0/2       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/2     Partner    No   Yes    No   No   No   Yes     Fast     Passive

    LACP protocol:        Receive State   Transmit State      Mux State
      ge-1/0/1                CURRENT      Fast periodic      Collecting
distributing
      ge-1/0/2                CURRENT      Fast periodic      Collecting
distributing
      ge-1/0/1 (active)       CURRENT      Fast periodic      Collecting
distributing
      ge-1/0/2 (standby)      CURRENT      Fast periodic      WAITING
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx    Illegal Rx
      ge-1/0/1                   0           0            0    0
      ge-1/0/2                   0           0            0    0
```

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
```

Copyright © 2015, Juniper Networks, Inc.

show lldp

Syntax

show lldp
<detail>

Release Information

Command introduced in Junos OS Release 9.6.

Description

On MX Series and T Series routers, 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 1328
show lldp detail on page 1328

Output Fields

Table 93 on page 1326 describes the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

Table 93: 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.
Interface	List of LLDP interfaces, showing status (**Enabled** or **Disabled**) and **Neighbor count** (**detail** only). 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*.
LLDP basic TLVs supported	List of basic LLDP TLVs supported by this device (**detail** only).

 Copyright © 2015, Juniper Networks, Inc.

Table 93: show lldp Output Fields *(continued)*

Field Name	Field Description
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 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
ge-0/0/0       Enabled
ge-0/0/1       Enabled
ge-0/0/4       Enabled
```

Sample Output

show lldp detail

```
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-0/0/0       Enabled    0
ge-0/0/1       Enabled    0
ge-0/0/4       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.
```


Copyright © 2015, Juniper Networks, Inc.

show lldp local-information

Syntax show lldp local-information

Release Information Command introduced in Junos OS Release 9.6.

Description On MX Series and T Series routers, 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 1331
show lldp local-information(Management Information Address Subtype is IPv6) on page 1331

Output Fields Table 94 on page 1329 describes the output fields for the **show lldp local-information** command. Output fields are listed in the approximate order in which they appear.

Table 94: 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 94: 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**.

Copyright © 2015, Juniper Networks, Inc.

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
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

On MX Series and T Series routers, 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*.

Options

interface ***interface-name***—(Optional) Display the neighbor information about a particular physical interface.

Required Privilege Level

view

Related Documentation

- clear lldp neighbor on page 926

List of Sample Output

show lldp neighbors on page 1334
show lldp neighbors interface ge-0/0/4 (Management Address is IPv4) on page 1334
show lldp neighbors interface ge-0/0/4 (Management Address is IPv6) on page 1335

Output Fields

Table 95 on page 1332 describes the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 95: show lldp neighbors Output Fields

Field Name	Field Description
LLDP Remote Devices Information	Information about remote devices.
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).

Copyright © 2015, Juniper Networks, Inc.

Table 95: show lldp neighbors Output Fields *(continued)*

Field Name	Field Description
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).
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).

Copyright © 2015, Juniper Networks, Inc.

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)
```


Copyright © 2015, Juniper Networks, Inc.

```
       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
```


Copyright © 2015, Juniper Networks, Inc.

show lldp remote-global-statistics

Syntax show lldp remote-global-statistics

Release Information Command introduced in Junos OS Release 9.6.

Description On MX Series and T Series routers, 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 1338

Output Fields Table 96 on page 1337 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 96: 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.

Copyright © 2015, Juniper Networks, Inc.

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
```


Copyright © 2015, Juniper Networks, Inc.

show lldp statistics

Syntax

show lldp statistics
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.6.

Description

On MX Series and T Series routers, display information about Link Layer Discovery Protocol (LLDP) statistics.

Options

interface ***interface-name***—(Optional) Display the statistics about a particular physical interface.

Required Privilege Level

view

Related Documentation

- clear lldp statistics on page 927

List of Sample Output

show lldp statistics on page 1340
show lldp statistics interface ge-0/1/1 on page 1340

Output Fields

Table 97 on page 1339 describes the output fields for the **show lldp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 97: 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*.
Received	Number of LLDP frames received on this interface.
Transmitted	Number of LLDP frames sent on this interface.
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.

Sample Output

show lldp statistics

```
user@host> show lldp statistics
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
--------- -------- ----------- ------------ ----------- ---------
ge-0/1/1  544      540         0            0           0
ge-0/1/2  540      500         0            0           0
ge-0/1/3  544      540         0            0           0
ge-0/1/4  544      540         0            0           0
ge-0/1/5  544      540         0            0           0
ge-0/1/6  544      540         0            0           0
ge-0/1/7  0        0           0            0           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
```


Copyright © 2015, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management delay-statistics

Syntax

show oam ethernet connectivity-fault-management delay-statistics
<**count** *entry-count*>
<**local-mep** *local-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 ***local-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*
- *clear oam ethernet connectivity-fault-management delay-statistics*
- show oam ethernet connectivity-fault-management interfaces on page 1349
- show oam ethernet connectivity-fault-management mep-database on page 1366
- show oam ethernet connectivity-fault-management mep-statistics on page 1377

List of Sample Output

show oam ethernet connectivity-fault-management delay-statistics on page 1343
show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 1343

Output Fields

Table 98 on page 1342 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 98: 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.

 Copyright © 2015, Juniper Networks, Inc.

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
```

Copyright © 2015, Juniper Networks, Inc.

```
  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
```

 Copyright © 2015, Juniper Networks, Inc.

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 1346
show oam ethernet connectivity-fault-management forwarding-state interface on page 1346
show oam ethernet connectivity-fault-management forwarding-state interface detail on page 1347
show oam ethernet connectivity-fault-management forwarding-state interfaceinterface-name on page 1348

Output Fields

Table 99 on page 1345 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 99: 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**
Level	Maintenance domain (MD) level.	**detail**

Copyright © 2015, Juniper Networks, Inc.

Table 99: show oam ethernet connectivity-fault-management forwarding-state Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```


Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

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*
- show oam ethernet connectivity-fault-management delay-statistics on page 1341
- show oam ethernet connectivity-fault-management mep-database on page 1366
- show oam ethernet connectivity-fault-management mep-statistics on page 1377

List of Sample Output

show oam ethernet connectivity-fault-management interfaces on page 1354
show oam ethernet connectivity-fault-management interfaces detail on page 1354
show oam ethernet connectivity-fault-management interfaces detail (One-Way ETH-DM) on page 1355
show oam ethernet connectivity-fault-management interfaces detail(Connection Protection TLV Configured) on page 1356
show oam ethernet connectivity-fault-management interfacesextensive on page 1357
show oam ethernet connectivity-fault-management interfaces level on page 1358

Copyright © 2015, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management interfaces (trunk ports) on page 1358

Output Fields Table 100 on page 1350 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 100: 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**
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**
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**

 Copyright © 2015, Juniper Networks, Inc.

Table 100: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**

Copyright © 2015, Juniper Networks, Inc.

Table 100: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**

 Copyright © 2015, Juniper Networks, Inc.

Table 100: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**

Copyright © 2015, Juniper Networks, Inc.

Table 100: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```


Copyright © 2015, Juniper Networks, Inc.

```
MEP status: running
Defects:
  Remote MEP not receiving CCM               : no
  Erroneous CCM received                     : yes
  Cross-connect CCM received                 : no
  RDI sent by some MEP                       : yes
Statistics:
  CCMs sent                                  : 76
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request        : 0
  1DMs sent                                  : 0
  Valid 1DMs received                        : 0
  Invalid 1DMs received                      : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                        : 0
  Invalid DMRs received                      : 0
  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: 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         : 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 interfaces detail (Connection Protection TLV Configured)

user@host**show 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         : 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
```


Copyright © 2015, Juniper Networks, Inc.

```
      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
    Statistics:
      CCMs sent                                     : 76
      CCMs received out of sequence                 : 0
      LBMs sent                                     : 0
      Valid in-order LBRs received                  : 0
      Valid out-of-order LBRs received              : 0
      LBRs received with corrupted data             : 0
      LBRs sent                                     : 0
      LTMs sent                                     : 0
      LTMs received                                 : 0
      LTRs sent                                     : 0
      LTRs received                                 : 0
      Sequence number of next LTM request           : 0
      1DMs sent                                     : 0
      Valid 1DMs received                           : 0
      Invalid 1DMs received                         : 0
      DMMs sent                                     : 0
      DMRs sent                                     : 0
      Valid DMRs received                           : 0
      Invalid DMRs received                         : 0
      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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

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 1360
show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 1360

Output Fields Table 101 on page 1359 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 101: 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.
Hop	Sequential hop count of the linktrace path.

Copyright © 2015, Juniper Networks, Inc.

Table 101: show oam ethernet connectivity-fault-management linktrace path-database Output Fields *(continued)*

Field Name	Field Description
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
```


Copyright © 2015, Juniper Networks, Inc.

```
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 1377

Output Fields

Table 102 on page 1362 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 102: 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.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-LM entry in the CFM database.

Copyright © 2015, Juniper Networks, Inc.

Table 102: show oam ethernet connectivity-fault-management loss-statistics Output Fields *(continued)*

Output Field Name	Field Description
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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%
```

Copyright © 2015, Juniper Networks, Inc.

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*
- show oam ethernet connectivity-fault-management delay-statistics on page 1341
- show oam ethernet connectivity-fault-management interfaces on page 1349
- show oam ethernet connectivity-fault-management mep-statistics on page 1377

List of Sample Output

show oam ethernet connectivity-fault-management mep-database on page 1371
show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM) on page 1372
show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 1373

Copyright © 2015, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 1373
show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured) on page 1373
show oam ethernet connectivity-fault-management mep-database on page 1374
show oam ethernet connectivity-fault-management mep-database (enhanced continuity measurement) on page 1375

Output Fields Table 103 on page 1367 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 103: 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**.
FRR Flag	LSR/LER forwarding the CCM Frame into a bypass tunnel is set. Its value can be **yes** or **no**.

Copyright © 2015, Juniper Networks, Inc.

Table 103: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
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**.
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**.

Copyright © 2015, Juniper Networks, Inc.

Table 103: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
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.
Sequence number of next LTM request	Sequence number of the next linktrace message request to be transmitted.

Copyright © 2015, Juniper Networks, Inc.

Table 103: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
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.
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.

 Copyright © 2015, Juniper Networks, Inc.

Table 103: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
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:
```

Copyright © 2015, Juniper Networks, Inc.

```
    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
```

 Copyright © 2015, Juniper Networks, Inc.

```
  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.

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

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*
- show oam ethernet connectivity-fault-management delay-statistics on page 1341
- show oam ethernet connectivity-fault-management interfaces on page 1349
- show oam ethernet connectivity-fault-management mep-database on page 1366

List of Sample Output

show oam ethernet connectivity-fault-management mep-statistics (CIR counters only) on page 1379
show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled) on page 1381
show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only) on page 1382

Copyright © 2015, Juniper Networks, Inc.

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled) on page 1383
show oam ethernet connectivity-fault-management mep-statistics on page 1385
show oam ethernet connectivity-fault-management mep-statistics remote-mep on page 1386

Output Fields Table 104 on page 1378 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 104: 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.

 Copyright © 2015, Juniper Networks, Inc.

Table 104: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields *(continued)*

Output Field Name	Field Description
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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%
```

 Copyright © 2015, Juniper Networks, Inc.

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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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 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 1389

Output Fields

Table 105 on page 1388 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 105: 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.
Maintenance association name	Maintenance association name.

Copyright © 2015, Juniper Networks, Inc.

Table 105: show oam ethernet connectivity-fault-management path-database Output Fields *(continued)*

Field Name	Field Description
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 931

List of Sample Output show oam ethernet connectivity-fault-management policer on page 1391
show oam ethernet connectivity-fault-management policer maintenance-domain md-name on page 1391
show oam ethernet connectivity-fault-management policer maintenance-domain md-name maintenance-association ma-name on page 1392

Output Fields Table 106 on page 1390 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 106: 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.
Level	Displays the maintenance domain level configured.

Copyright © 2015, Juniper Networks, Inc.

Table 106: show oam ethernet connectivity-fault-management policer Output Fields *(continued)*

Field Name	Field Description
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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

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 1396
show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) on page 1396

Output Fields

Table 107 on page 1393 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 107: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields

Output Field Name	Output Field Description
Maintenance domain	Name of the maintenance domain.

Copyright © 2015, Juniper Networks, Inc.

Table 107: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields *(continued)*

Output Field Name	Output Field Description
Level	Level of the maintenance domain level configured.
Maintenance association	Name of the maintenance association.
Local MEP id	Numeric identifier of the local MEP.
Remote MEP id	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
Delay weight	Calculation weight of delay.
Delay variation weight	Calculation weight of delay variation.
DMM sent	Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.
DMM skipped for threshold hit	Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.
DMM skipped for threshold hit window	Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
DMR received	Number of delay measurement reply (DMR) frames received.
DMR out of sequence	Total number of DMR out of sequence packets received.
DMR received with invalid time stamps	Total number of DMR frames received with invalid timestamps.

Copyright © 2015, Juniper Networks, Inc.

Table 107: 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.

Copyright © 2015, Juniper Networks, Inc.

Table 107: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields *(continued)*

Output Field Name	Output Field Description
Last Received SLR frame TXFCb(t	Value of the local counter 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 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
```


Copyright © 2015, Juniper Networks, Inc.

```
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*
- clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 932
- show oam ethernet connectivity-fault-management interfaces on page 1349
- show oam ethernet connectivity-fault-management mep-database on page 1366
- show oam ethernet connectivity-fault-management mep-statistics on page 1377

List of Sample Output

show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 1399

Output Fields

Table 108 on page 1399 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.

 Copyright © 2015, Juniper Networks, Inc.

Table 108: 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

Copyright © 2015, Juniper Networks, Inc.

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 %)
```


Copyright © 2015, Juniper Networks, Inc.

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 109 on page 1401 lists the output fields for the **show oam ethernet evc** command. Output fields are listed in the approximate order in which they appear.

Table 109: 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
```

Copyright © 2015, Juniper Networks, Inc.

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 1402

Output Fields

Table 110 on page 1402 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 110: 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
```


Copyright © 2015, Juniper Networks, Inc.

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 1404

Output Fields

Table 111 on page 1403 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 111: 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.
Disabled interfaces	Name of the interfaces that are disabled.

Copyright © 2015, Juniper Networks, Inc.

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
```


Copyright © 2015, Juniper Networks, Inc.

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 1405

Output Fields

Table 112 on page 1405 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 112: 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
```

Copyright © 2015, Juniper Networks, Inc.

```
ge-0/0/0.1278        100      0.f0:ff:ff:ff:ff:ff                 0
```


Copyright © 2015, Juniper Networks, Inc.

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 1411
show oam ethernet link-fault-management detail on page 1411

Output Fields

Table 113 on page 1407 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 113: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Indicates the status of the established link. • **Fail**—A link fault condition exists. • **Running**—A link fault condition does not exist. • **ISSU**—The local end is in ISSU.	All levels
Discovery state	State of the discovery mechanism: • **Passive Wait** • **Send Any** • **Send Local Remote** • **Send Local Remote Ok** • **Fault**	All levels
Peer address	Address of the OAM peer.	All levels

Copyright © 2015, Juniper Networks, Inc.

Table 113: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the interface. Possible values are described in the “Link Flags” section under *Common Output Fields Description*. • **Remote-Stable**—Indicates remote OAM client acknowledgment of and satisfaction with local OAM state information. **False** indicates that remote DTE either has not seen or is unsatisfied with local state information.**True** indicates that remote DTE has seen and is satisfied with local state information. • **Local-Stable**—Indicates local OAM client acknowledgment of and satisfaction with remote OAM state information. **False** indicates that local DTE either has not seen or is unsatisfied with remote state information. **True** indicates that local DTE has seen and is satisfied with remote state information. • **Remote-State-Valid**—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. **False** indicates that OAM client has not seen remote state information. **True** indicates that the OAM client has seen remote state information.	All levels
Remote loopback status	Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
Remote entity information	Remote entity information. • **Remote MUX action**—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • **Remote parser action**—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • **Discovery mode**—Indicates whether discovery mode is active or inactive. • **Unidirectional mode**—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes. • **Remote loopback mode**—Indicates whether remote loopback is supported or unsupported. • **Link events**—Indicates whether interpreting link events is supported or unsupported on the remote peer. • **Variable requests**—Indicates whether variable requests are supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer. • Also indicates if the remote end is in ISSU.	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**

 Copyright © 2015, Juniper Networks, Inc.

Table 113: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Loopback control	The total number of loopback control PDUs received.	**detail**
Organization specific	The total number of vendor organization specific PDUs received.	**detail**
OAM Transmit Statistics		
Information	The total number of information PDUs transmitted.	**detail**
Event	The total number of event notification PDUs transmitted.	**detail**
Variable request	The total number of variable request PDUs transmitted.	**detail**
Variable response	The total number of variable response PDUs transmitted.	**detail**
Loopback control	The total number of loopback control PDUs transmitted.	**detail**
Organization specific	The total number of vendor organization specific PDUs transmitted.	**detail**
OAM Received Symbol Error Event information		
Events	The number of symbol error event TLVs that have been received since the OAM sublayer was reset.	**detail**
Window	The symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	**detail**
Threshold	The number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	The number of symbol errors in the period reported in the received event PDU.	**detail**
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**

Copyright © 2015, Juniper Networks, Inc.

Table 113: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Total errors	The number of errored frames that have been reported in received event TLVs since the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	**detail**
OAM Received Frame Period Error Event Information		
Events	The number of frame seconds errors event TLVs that have been received since the OAM sublayer was reset.	**detail**
Window	The duration of the frame seconds window.	**detail**
Threshold	The number of frame seconds errors in the period.	**detail**
Errors in period	The number of frame seconds errors in the period.	**detail**
Total errors	The number of frame seconds errors that have been reported in received event TLVs since the OAM sublayer was reset.	**detail**
OAM Transmitted Symbol Error Event Information		
Events	The number of symbol error event TLVs that have been transmitted since the OAM sublayer was reset.	**detail**
Window	The symbol error event window in the transmitted PDU.	**detail**
Threshold	The number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	The number of symbol errors in the period reported in the transmitted event PDU.	**detail**
Total errors	The number of errored symbols reported in event TLVs that have been transmitted since the OAM sublayer was reset.	**detail**
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		

 Copyright © 2015, Juniper Networks, Inc.

Table 113: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
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 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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

 Copyright © 2015, Juniper Networks, Inc.

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 114 on page 1413 lists the output fields for the **show oam ethernet lmi** command. Output fields are listed in the approximate order in which they appear.

Table 114: 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.
EVC Indentifier	EVC name.

Table 114: show oam ethernet lmi Output Fields *(continued)*

Field Name	Field Description
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
```


Copyright © 2015, Juniper Networks, Inc.

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 1415

Output Fields Table 115 on page 1415 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 115: 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**

Copyright © 2015, Juniper Networks, Inc.

```
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
```


Copyright © 2015, Juniper Networks, Inc.

show pppoe interfaces

Syntax

show pppoe interfaces
<brief | detail | extensive>
<pp0.*logical*>

Release Information

Command introduced before Junos OS Release 7.4.

Description

(J Series Services Routers, M120 routers, M320 routers, and MX Series routers only) Display session-specific information about PPPoE interfaces.

Options

none—Display interface information for all PPPoE interfaces.

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

extensive—(J Series Services Routers) (Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.

pp0.*logical*—(Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 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 1419
show pppoe interfaces (Status for the Specified Interface) on page 1419
show pppoe interfaces brief on page 1420
show pppoe interfaces detail on page 1420
show pppoe interfaces extensive (J Series Services Routers only) on page 1420
show pppoe interfaces (PPPoE Subscriber Interface with ACI Interface Set) on page 1420

Output Fields

Table 116 on page 1417 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 116: 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

Copyright © 2015, Juniper Networks, Inc.

Table 116: show pppoe interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Session ID	Session ID.	All levels
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
Auto-reconnect timeout	(J Series Services Routers only) Time after which to try to reconnect after a PPPoE session is terminated, in seconds.	**detail extensive** none
Idle timeout	(J Series Services Routers only) Length of time (in seconds) that a connection can be idle before disconnecting.	**detail extensive** none
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

 Copyright © 2015, Juniper Networks, Inc.

Table 116: 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**
Timeout	(J Series Services Routers only) Information about timeouts that occurred during the PPPoE session: • **PADI**—No PADO packet has been received within the timeout period. • **PADO**—No PADR packet has been received within the timeout period. (This value is always zero and is not supported.) • **PADR**—No PADS packet has been received within the timeout period.	**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: 00:00:00:00:00:00,
  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: 00:90:1A:42:0A: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
```

Copyright © 2015, Juniper Networks, Inc.

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         00:90:1A:42:0A:C1
pp0.1             ge-2/0/3.2          Session Up  28         00:90:1A:42:0A:C1
pp0.1073741824    ge-2/0/3.1          Session Up  29         00:90:1A:42:0A:C1
pp0.1073741825    ge-2/0/3.1          Session Up  30         00:90:1A:42:0A:C1
pp0.1073741826    ge-2/0/3.1          Session Up  31         00:90:1A:42:0A: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: 00:00:00:00:00:00,
  Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
  Underlying interface: at-5/0/0.0 Index 71
```

show pppoe interfaces extensive (J Series Services Routers only)

```
user@host> show pppoe interfaces pp0.1 extensive
pp0.1 Index 66
  State: Down, Session ID: 26, Type: Static,
  Service name: None, Configured AC name: sapphire,
  Session AC name: None, Remote MAC address: 00:00:00:00:00:00,
  Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
  Underlying interface: ge-3/0/3.1 Index 71
  PacketType                  Sent          Received
    PADI                         0                 0
    PADO                         0                 0
    PADR                         0                 6
    PADS                         6                 0
    PADT                         6                 0
    Service name error           0                 0
    AC system error              0                 0
    Generic error                0                 0
    Malformed packets            0                 0
    Unknown packets              0                 0
  Timeout
    PADI                         0
    PADO                         0
    PADR                         0
```

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: 00:00:64:39:01:02,
  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
```

 Copyright © 2015, Juniper Networks, Inc.

show pppoe service-name-tables

Syntax show pppoe service-name-tables
<*table-name*>

Release Information Command introduced in Junos OS Release 10.0.

Description (M120 routers, M320 routers, and MX Series routers only) 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 220
- *Verifying and Managing Dynamic PPPoE Configuration*

List of Sample Output show pppoe service-name-tables on page 1422
show pppoe service-name-tables (For the Specified Table Name) on page 1422

Output Fields Table 117 on page 1421 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 117: 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
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

Copyright © 2015, Juniper Networks, Inc.

Table 117: show pppoe service-name-tables Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
```

 Copyright © 2015, Juniper Networks, Inc.

```
      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

(M120 routers, M320 routers, and MX Series routers only) 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 220
- *Verifying and Managing Dynamic PPPoE Configuration*

List of Sample Output

show pppoe sessions (For All Active Sessions) on page 1425
show pppoe sessions (For All Active Sessions Matching the Agent Circuit Identifier) on page 1425

Output Fields

Table 118 on page 1424 lists the output fields for the **show pppoe sessions** command. Output fields are listed in the approximate order in which they appear.

Table 118: 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
Underlying interface	Interface on which PPPoE is running.	none

Copyright © 2015, Juniper Networks, Inc.

Table 118: show pppoe sessions Output Fields *(continued)*

Field Name	Field Description	Level of Output
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         00:90:1A:42:0A:C1
pp0.1             ge-2/0/3.2          Session Up  28         00:90:1A:42:0A:C1
pp0.1073741824    ge-2/0/3.1          Session Up  29         00:90:1A:42:0A:C1
pp0.1073741825    ge-2/0/3.1          Session Up  30         00:90:1A:42:0A:C1
pp0.1073741826    ge-2/0/3.1          Session Up  31         00:90:1A:42:0A: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         00:90:1A:42:0A:C1
pp0.1             ge-2/0/3.2          Session Up  28         00:90:1A:42:0A:C1
```

Copyright © 2015, Juniper Networks, Inc.

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

(J Series Services Routers, M120 routers, M320 routers, and MX Series routers only) 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. Supported for M120 routers, M320 routers, and MX Series routers only.

Required Privilege Level

view

Related Documentation

- *show ppp address-pool*
- show pppoe underlying-interfaces on page 1428

List of Sample Output

show pppoe statistics on page 1427
show pppoe statistics (For the Specified Underlying Interface Only) on page 1427

Output Fields

Table 119 on page 1426 lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 119: 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.

Copyright © 2015, Juniper Networks, Inc.

Table 119: 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
```

Copyright © 2015, Juniper Networks, Inc.

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

(M120, M320, and MX Series routers only) 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*

List of Sample Output

show pppoe underlying-interfaces brief on page 1430
show pppoe underlying-interfaces detail on page 1431
show pppoe underlying-interfaces extensive on page 1431
show pppoe underlying-interfaces extensive (PPPoE client in lockout condition) on page 1432
show pppoe underlying-interfaces lockout on page 1432
show pppoe underlying-interfaces detail (Autosensing Configured for ACI-based Dynamic VLANs) on page 1432

Output Fields

Table 120 on page 1428 lists the output fields for the **show pppoe underlying-interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 120: 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

Copyright © 2015, Juniper Networks, Inc.

Table 120: show pppoe underlying-interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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 to enable creation of (autosense) 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.	**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.	**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

Table 120: show pppoe underlying-interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
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.	**extensive**
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**zero) 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**

Copyright © 2015, Juniper Networks, Inc.

```
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
    00:10:94:00:00:01               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,
```


Copyright © 2015, Juniper Networks, Inc.

```
Agent Circuit Identifier: Autosensing,
Service Name Table: None,
Duplicate Protection: On, Short Cycle Protection: Off,
Direct Connect: Off,
AC Name: nbc,
Short Cycle Protection Level: circuit-id,
```

Copyright © 2015, Juniper Networks, Inc.

show pppoe version

Syntax show pppoe version

Release Information Command introduced before Junos OS Release 7.4.

Description (J Series Services Routers, 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 1434

Output Fields Table 121 on page 1434 lists the output fields for the **show pppoe version** command. Output fields are listed in the approximate order in which they appear.

Table 121: 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
```

 Copyright © 2015, Juniper Networks, Inc.

```
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.

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 1438
- show protection-group ethernet-ring interface on page 1440
- show protection-group ethernet-ring node-state on page 1443
- show protection-group ethernet-ring statistics on page 1446
- show protection-group ethernet-ring vlan on page 1449

List of Sample Output

show protection-group ethernet-ring aps (EX Switches) on page 1437
show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers) on page 1437
show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers) on page 1437
show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers) on page 1437
show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers) on page 1437

Output Fields Table 122 on page 1436 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 122: show protection-group ethernet-ring aps Output Fields

Field Name	Field Description
Ethernet Ring Name	Name configured for the Ethernet ring.
Request/State	Status of the Ethernet ring RAPS messages. • **NR**—Indicates there is no request for APS on the ring. • **SF**—Indicates there is a signal failure on the ring.
No Flush	State of the ring flushing: **No** (normal) or **Yes** (failure).
Ring Protection Link Blocked	Blocking on the ring protection link: **Yes** or **No**.

 Copyright © 2015, Juniper Networks, Inc.

Table 122: show protection-group ethernet-ring aps Output Fields *(continued)*

Field Name	Field Description
Originator	Whether this node is the ring originator: **Yes** or **No**.
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 MX Routers)

```
user@host> 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
```

show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers)

```
user@host> 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
```

show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers)

```
user@host> 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
```

show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               SF             No        Yes

Originator  Remote Node ID
Yes         00:00:00:00:00:00
```

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.

Description

On MX Series routers, display data channel information for all Ethernet ring protection groups or for a specific Ethernet ring protection group.

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 1436
- show protection-group ethernet-ring interface on page 1440
- show protection-group ethernet-ring node-state on page 1443
- show protection-group ethernet-ring statistics on page 1446
- show protection-group ethernet-ring vlan on page 1449

List of Sample Output

show protection-group ethernet-ring data-channel on page 1439

Output Fields

Table 123 on page 1438 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 123: show protection-group ethernet-ring data-channel Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet ring.
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.

 Copyright © 2015, Juniper Networks, Inc.

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
ge-1/0/3     71               forwarding
ge-1/0/4     82               forwarding

Ethernet ring data channel information for protection group pg302
Interface    STP index        Forward State
ge-1/0/3     52               forwarding
ge-1/0/4     91               forwarding
```

Copyright © 2015, Juniper Networks, Inc.

show protection-group ethernet-ring interface

Syntax **show protection-group ethernet-ring interface**

Release Information Command introduced in Junos OS Release 9.4.

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 1438
- show protection-group ethernet-ring aps on page 1436
- show protection-group ethernet-ring node-state on page 1443
- show protection-group ethernet-ring statistics on page 1446
- show protection-group ethernet-ring vlan on page 1449

List of Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node) on page 1441
show protection-group ethernet-ring interface (Owner Node MX Series Router) on page 1441
show protection-group ethernet-ring interface (EX Series Switch Ring Node) on page 1441
show protection-group ethernet-ring interface (MX Series Router Ring Node) on page 1441

Output Fields Table 124 on page 1440 lists the output fields for both the EX Series switch and the MX Series router **show protection-group ethernet-ring interface** commands. Output fields are listed in the approximate order in which they appear.

Table 124: 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.
Control Channel	(MX Series router only) Logical unit configured on the physical interface. • **NR**—Indicates there is no request for APS on the ring. • **SF**—Indicates there is a signal failure on the ring.
Forward State	State of the ring forwarding on the interface: **discarding** or **forwarding**.

 Copyright © 2015, Juniper Networks, Inc.

Table 124: MX Series Routers show protection-group ethernet-ring interface Output Fields *(continued)*

Field Name	Field Description
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  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
```

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 (MX Series Router Ring Node)

```
user@host> 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
```

Copyright © 2015, Juniper Networks, Inc.

```
Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```


Copyright © 2015, Juniper Networks, Inc.

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.
Command introduced in Junos OS Release 12.1 for EX Series 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 1438
- show protection-group ethernet-ring aps on page 1436
- show protection-group ethernet-ring interface on page 1440
- show protection-group ethernet-ring statistics on page 1446
- show protection-group ethernet-ring vlan on page 1449

List of Sample Output

show protection-group ethernet-ring node-state (EX Series Switch) on page 1444
show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation) on page 1444
show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation) on page 1444
show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition) on page 1445
show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition) on page 1445
show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion) on page 1445

Output Fields Table 125 on page 1443 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 125: show protection-group ethernet-ring node-state Output Fields

Field Name	Field Description
Ring Name/Ethernet Ring	Name configured for the Ethernet ring.

Table 125: 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 no protection-switching request active or pending in the ring. When the ring is in idle state, it is blocked at RPL link. • **protected**—Indicates that there is a protection switch on the ring due to signal failure condition on the ring link.
Event	Events on the ring. • **NR-RB**—Indicates there is no APS request and the ring link is blocked on the ring owner node. • **NR**—Indicates there is no APS request pending in the ring. • **Local SF**—Indicates there is signal failure on one or both the ring links of the node. • **Remote SF**—Indicates there is signal failure on ring links of any other node of the ring.
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**.
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 (EX Series Switch)

```
user@switch> show protection-group ethernet-ring node-state
Ring Name APS State Event RPL Owner WTR Timer Guard Timer Op State
erp1      idle       NR-RB   yes      disabled    disabled    operational
```

show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         RPL Owner
pg101             idle          NR-RB         Yes

Restore Timer  Guard Timer  Operation state
disabled       disabled      operational
```

show protection-group ethernet-ring node-state (MX Series Router - 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
```


Copyright © 2015, Juniper Networks, Inc.

```
Restore Timer  Guard Timer  Operation state
disabled       disabled     operational
```

show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         RPL Owner
pg101             protected     remote SF           Yes

Restore Timer  Guard Timer  Operation state
disabled       disabled     operational
```

show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         RPL Owner
pg102             protected     local SF       No

Restore Timer  Guard Timer  Operation state
disabled       disabled     operational
```

show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : pg_major
APS State                       : protected
Event                           : WTR running (time to expire: 269 sec)
Ring Protection Link Owner      : Yes
Restore Timer                   : running
Guard Timer                     : disabled
Operation state                 : operational
```

show protection-group ethernet-ring statistics

Syntax **show protection-group ethernet-ring statistics <group-name *group-name*>**

Release Information Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.

Options **group-name**—Protection group for which to display statistics. In 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 data-channel on page 1438
- show protection-group ethernet-ring aps on page 1436
- show protection-group ethernet-ring node-state on page 1443
- show protection-group ethernet-ring interface on page 1440
- show protection-group ethernet-ring vlan on page 1449

List of Sample Output

show protection-group ethernet-ring statistics (EX Switch) on page 1447
show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router) on page 1447
show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router) on page 1447
show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router) on page 1447
show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on MX Router) on page 1448

Output Fields Table 126 on page 1446 lists the output fields for the **show protection-group ethernet-ring statistics** command. Output fields are listed in the approximate order in which they appear.

Table 126: 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 sent	Number of Ring Automatic Protection Switching (RAPS) messages sent. (On MX Series switches only)
RAPS received	Number of RAPS messages received. (On MX Series switches only)

Copyright © 2015, Juniper Networks, Inc.

Table 126: show protection-group ethernet-ring statistics Output Fields *(continued)*

Field Name	Field Description
Local SF	Number of times a signal failure (SF) has occurred locally.
Remote SF	Number of times a signal failure (SF) has occurred anywhere else on the ring.
NR event	Number of times a No Request (NR) event has occurred on the ring.
NR-RB event	Number of times a No Request, Ring Blocked (NR-RB) event has occurred on the ring.

Sample Output

show protection-group ethernet-ring statistics (EX 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 (Owner Node, Normal Operation on 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
```

show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router)

```
user@host> 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
```

show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router)

```
user@host> 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
```

Copyright © 2015, Juniper Networks, Inc.

show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on MX Router)

```
user@host> 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
```

 Copyright © 2015, Juniper Networks, Inc.

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.

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 1436
- show protection-group ethernet-ring data-channel on page 1438
- show protection-group ethernet-ring interface on page 1440
- show protection-group ethernet-ring node-state on page 1443
- show protection-group ethernet-ring statistics on page 1446

List of Sample Output

show protection-group ethernet-ring vlan on page 1450
show protection-group ethernet-ring vlan brief on page 1450
show protection-group ethernet-ring vlan detail on page 1450
show protection-group ethernet-ring vlan group-name vkm01 on page 1450

Output Fields

Table 127 on page 1449 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 127: 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.
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.

Table 127: show protection-group ethernet-ring vlan Output Fields *(continued)*

Field Name	Field Description
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
ge-2/0/8      100     130         default-switch/bd100
ge-2/0/4      100     126         default-switch/bd100
```

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
ge-2/0/8      100     130         default-switch/bd100
ge-2/0/4      100     126         default-switch/bd100
```

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                : ge-2/0/8
Vlan                          : 100
STP index                     : 130
Bridge Domain                 : default-switch/bd100

Interface name                : ge-2/0/4
Vlan                          : 100
STP index                     : 126
Bridge Domain                 : default-switch/bd100
```

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
ge-2/0/8      100     130         default-switch/bd100
ge-2/0/4      100     126         default-switch/bd100
```


Copyright © 2015, Juniper Networks, Inc.

show vrrp

Syntax

show vrrp
<brief | detail | extensive | summary>
<interface *interface-name* <group *number*>>
<logical-system (all | *logical-system-name*)>
<nsr>
<track <interfaces>>

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, group number.

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

track <interfaces>—(Optional) Display information and status about VRRP track interfaces.

Required Privilege Level

view

Related Documentation

- *clear vrrp*

List of Sample Output

show vrrp on page 1456
show vrrp brief on page 1457
show vrrp detail (IPv6) on page 1457
show vrrp detail (Route Track) on page 1457
show vrrp extensive on page 1457
show vrrp interface on page 1458
show vrrp nsr on page 1459
show vrrp summary on page 1460
show vrrp track detail on page 1460
show vrrp track summary on page 1460

Output Fields

Table 128 on page 1452 lists the output fields for the **show vrrp** command. Output fields are listed in the approximate order in which they appear

Table 128: 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	Nonerrored 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**

 Copyright © 2015, Juniper Networks, Inc.

Table 128: 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 or **Int 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. • **transition**—The interface is changing between being the backup and being the master router.	**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**
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**

Copyright © 2015, Juniper Networks, Inc.

Table 128: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
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	Name of the tracked interface.	**detail extensive**
Int state/Interface 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**

Copyright © 2015, Juniper Networks, Inc.

Table 128: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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**

Table 128: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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 Dead 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
```


Copyright © 2015, Juniper Networks, Inc.

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 1456.

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
  Priority: 200, Advertisement interval: 1, Authentication type: none
  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-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
```

Copyright © 2015, Juniper Networks, Inc.

```
    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
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

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
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: fe80::12:1:1:99,
fec0::12:1:1:99
  Advertisement timer: 0.789s, Master router: fe80::12:1:1:1
  Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
  Virtual MAC: 00:00:5e:00:02:01
```


Copyright © 2015, Juniper Networks, Inc.

```
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent            :         205
    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
```

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
```

Copyright © 2015, Juniper Networks, Inc.

```
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
```

show vrrp track detail

```
user@host> show vrrp track detail
Tracked interface: ae1.211
  State: up, Speed: 400m
  Incurred priority cost: 0
  Threshold    Priority cost   Active
  400m         10
  300m         60
  200m         110
  100m         160
  down         190
  Tracking VRRP interface: ae0.210, Group: 1
    VR State: master
    Current priority: 200, Configured priority: 200
    Priority hold-time: disabled,    Remaining-time: 50.351
```

show vrrp track summary

```
user@host> show vrrp track summary
Track if      State      Speed       VRRP if    Group   VR State    Current priority
ae1.211       up          400m        ae0.210     1      master      200
```


Copyright © 2015, Juniper Networks, Inc.

traceroute ethernet

Syntax

traceroute ethernet (*mac-address* | *mep-id*)
maintenance-association *ma-name*
maintenance-domain *md-name*
ttl *value*
<wait *seconds*>

Release Information

Command introduced in Junos OS Release 9.0.
mep-id option introduced in Junos OS Release 9.1.

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

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 1462

Output Fields

Table 129 on page 1461 lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 129: 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).

Copyright © 2015, Juniper Networks, Inc.

Table 129: traceroute ethernet Output Fields *(continued)*

Field Name	Field Description
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.

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
```


Copyright © 2015, Juniper Networks, Inc.

PART 6

Index

Copyright © 2015, Juniper Networks, Inc.

Index

Symbols

A

 Copyright © 2015, Juniper Networks, Inc.

C

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

G

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

N

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.

Copyright © 2015, Juniper Networks, Inc.

 Copyright © 2015, Juniper Networks, Inc.