JUNIPER NETWORKS

Junos® OS

IEEE 802.1ag OAM Connectivity-Fault Management Feature Guide for Routing Devices

Release
14.1

Published: 2014-05-15

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos® OS IEEE 802.1ag OAM Connectivity-Fault Management Feature Guide for Routing Devices
14.1

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

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at 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:

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

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming

configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xiii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xiv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** *metric***>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** *community-ids* **]**

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

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

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 send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

CHAPTER 1

IEEE 802.1ag OAM Connectivity-Fault Management

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 might comprise one or more service instances. Junos OS supports IEEE 802.1ag connectivity fault management.

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 on Aggregated Ethernet Interfaces

In Junos OS Release 9.3 and later, CFM is supported on aggregated Ethernet interfaces. CFM sessions (on aggregated Ethernet interfaces) that operate in centralized mode are processed on the Routing Engine. However, in releases before Junos OS Release 13.3, graceful Routing Engine switchover (GRES) was not supported on aggregated Ethernet interfaces. Starting with Junos OS Release 13.3, CFM sessions can also operate in distributed mode, in which case they are processed on the Flexible PIC Concentrator (FPC). As a result, GRES is supported on aggregated Ethernet interfaces.

CFM sessions are distributed by default. To disable the distribution of CFM sessions on aggregated Ethernet interfaces and to 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. CFM sessions on aggregated Ethernet interfaces can operate in mixed mode only if the **no-aggregate-delegate-processing** statement is enabled. When this statement is enabled all aggregated Ethernet sessions operate in centralized mode.

NOTE: **You need to restart the connectivity-fault management process (cfmd) if the no-aggregate-delegate-processing knob is enabled after CFM is configured.**

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.

Connectivity Fault Management Key Elements

Figure 1 on page 5 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 1: Relationship Among MEPs, MIPs, and Maintenance Domain Levels

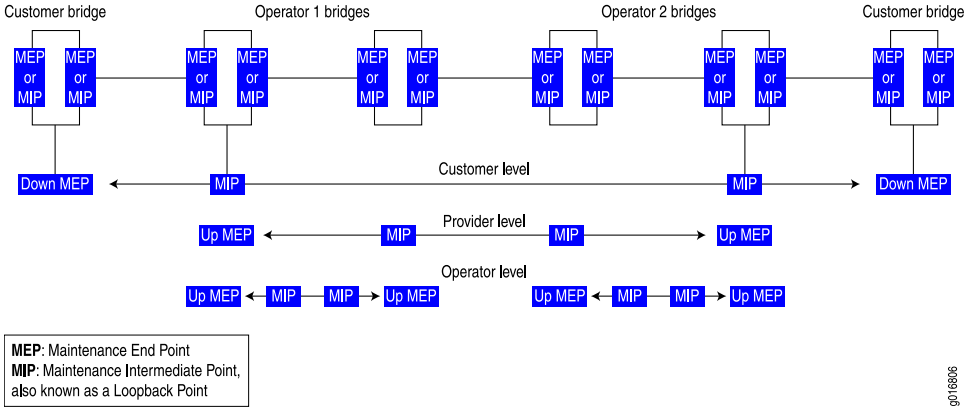

NOTE: Maintenance intermediate points (MIP) are not supported on the ACX Series routers.

A maintenance association is a set of MEPs configured with the same maintenance association identifier and maintenance domain level. Figure 2 on page 5 shows the hierarchical relationships between the Ethernet bridge, maintenance domains, maintenance associations, and MEPs.

Figure 2: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs

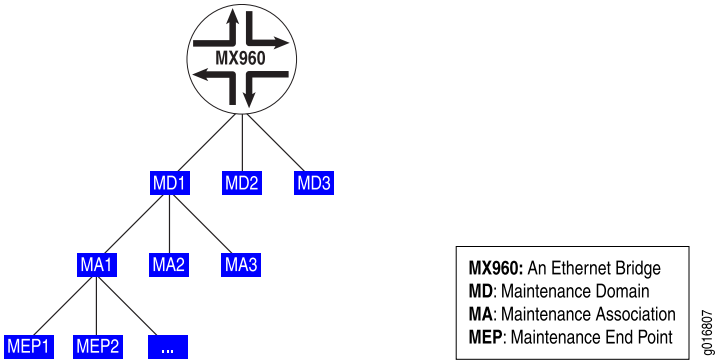

Best Practices for Configuring 802.1ag Ethernet OAM for VPLS

BEST PRACTICE: The logical interfaces in a VPLS routing instance might 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 3 and you should further observe the additional requirements described in this section.

For MX Series routers, the normalized VLAN is specified using one of the following configuration statements in the VPLS routing instance:

- vlan-id *vlan-number*
- vlan-id none
- vlan-tags outer *outer-vlan-number* inner *inner-vlan-number*

You must configure vlan-maps explicitly on all interfaces belonging to the routing instance.

The following forwarding path considerations must be observed:

- Packet receive path:
  - This is the forwarding path for packets received on the interfaces.
  - 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.
- Packet transmit path:
  - The JUNOS Software uses the router's hardware-based forwarding for CPU-generated packets.
  - For Down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - In MX series routers, for Up MEPs, the packet must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packet to be forwarded with this flood route.
  - The router also uses implicit-based forwarding for CPU generated packets. The result is for the flood next hop tied to the flood route to be tied to the filter term. The filter term uses match criteria to correctly identify the host- generated packets.

Related Documentation

- connectivity-fault-management on page 103
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15
- Continuity Check Protocol on page 16
- Configuring a Maintenance Endpoint on page 18
- Configuring a Connectivity Fault Management Action Profile on page 23

- Configuring Linktrace Protocol in CFM on page 26
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- Configuring Rate Limiting of Ethernet OAM Messages on page 52
- *Ethernet Interfaces*

Configuring CCM 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" on page 58.
- You can apply an action profile to provide faster convergence for dual-homed multipoint-to-multipoint network topologies. See "Configuring Faster Convergence" on page 59.
- 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 60.
- 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 Different ID" on page 61.

Related Documentation

- Configuring Faster Protection Switching on page 58
- Configuring Faster Convergence on page 59
- Configuring a Primary VLAN ID for Increased Flexibility on page 60
- Configuring a Remote Maintenance Association to Accept Different ID on page 61

PART 2

Configuration

CHAPTER 2

IEEE 802.1ag OAM Connectivity-Fault Management

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.

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

- Configuring the Maintenance Domain Name Format on page 12
- Configuring the Maintenance Domain Level on page 12

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

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 103
- maintenance-domain on page 119
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15
- Continuity Check Protocol on page 16

- Configuring a Maintenance Endpoint on page 18
- Configuring a Connectivity Fault Management Action Profile on page 23
- Configuring Linktrace Protocol in CFM on page 26
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- Configuring Rate Limiting of Ethernet OAM Messages on page 52
- *Ethernet Interfaces*

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 13
- Configuring the Maintenance Domain Bridge Domain on page 14
- Configuring the Maintenance Domain Instance on page 14
- Configuring the Maintenance Domain MIP Half Function on page 14

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 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 102
- connectivity-fault-management on page 103
- instance on page 112
- mip-half-function on page 121
- virtual-switch on page 133
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3

- Creating the Maintenance Domain on page 11
- Creating a Maintenance Association on page 15
- Continuity Check Protocol on page 16
- Configuring a Maintenance Endpoint on page 18
- Configuring a Connectivity Fault Management Action Profile on page 23
- Configuring Linktrace Protocol in CFM on page 26
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- Configuring Rate Limiting of Ethernet OAM Messages on page 52
- *Ethernet Interfaces*

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 103**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Continuity Check Protocol on page 16
- Configuring a Maintenance Endpoint on page 18

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

You can configure the following continuity check protocol parameters:

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

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.

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 103**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15
- Configuring a Maintenance Endpoint on page 18
- Configuring a Connectivity Fault Management Action Profile on page 23
- Configuring Linktrace Protocol in CFM on page 26
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- Configuring Rate Limiting of Ethernet OAM Messages on page 52
- *Managing Continuity Measurement Statistics*
- *Ethernet Interfaces*

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 on page 18
- Configuring the Maintenance Endpoint Direction on page 19
- Configuring the Maintenance Endpoint Interface on page 20
- Configuring the Maintenance Endpoint Priority on page 20
- Configuring the Maintenance Endpoint Lowest Priority Defect on page 20
- Configuring a Remote Maintenance Endpoint on page 21
- Configuring a Remote Maintenance Endpoint Action Profile on page 21
- Configuring Maintenance Endpoint Service Protection on page 22

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.

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.

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 3 on page 20 describes the available lowest priority defect options.

Table 3: 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.

Table 3: Lowest Priority Defect Options *(continued)*

Option	Description
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;
          }
        }
      }
    }
  }
}
```

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 4 on page 22 describes the available service protection options.

Table 4: Service Protection Options

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

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 103
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15
- Continuity Check Protocol on page 16
- Configuring a Connectivity Fault Management Action Profile on page 23
- Configuring Linktrace Protocol in CFM on page 26
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- Configuring Rate Limiting of Ethernet OAM Messages on page 52
- *Ethernet Interfaces*

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 24
- Configuring the Default Actions of a CFM Action Profile on page 24
- Configuring a CFM Action Profile Event on page 25

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 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 103**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Creating a Maintenance Association on page 15
- Configuring a Maintenance Endpoint on page 18
- *Ethernet Interfaces*

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 103**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15
- Continuity Check Protocol on page 16
- Configuring a Maintenance Endpoint on page 18
- Configuring a Connectivity Fault Management Action Profile on page 23
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- Configuring Rate Limiting of Ethernet OAM Messages on page 52
- *Ethernet Interfaces*

Configuring Ethernet Local Management Interface

- Ethernet Local Management Interface Overview on page 27
- Configuring the Ethernet Local Management Interface on page 29
- Example E-LMI Configuration on page 31

Ethernet Local Management Interface Overview

MX Series routers with Gigabit Ethernet (**ge**), 10-Gigabit Ethernet (**xe**), or Aggregated Ethernet (**ae**) interfaces support the Ethernet Local Management Interface (E-LMI). 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 3 on page 28.

Figure 3: Scope of the E-LMI Protocol

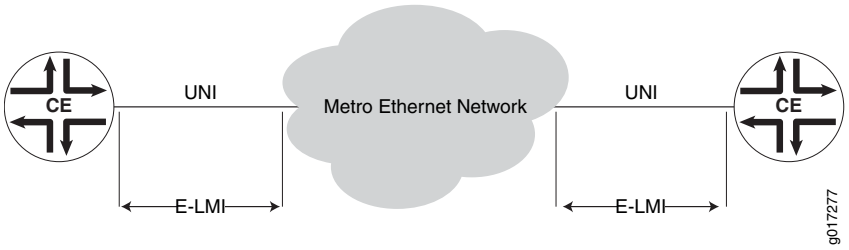

The E-LMI implementation on MX Series routers includes only the PE side of the E-LMI protocol.

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)
- 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 29
- Assigning the OAM Protocol to an EVC on page 29
- Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC on page 30

Configuring an OAM Protocol (CFM)

For information on configuring the OAM protocol (CFM), see "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 3.

Assigning the OAM Protocol to an EVC

To configure an EVC, you must specify a name for the EVC using the **evcs***evc-id* statement at the **[edit protocols oam ethernet]** hierarchy level. You can set the EVC protocol for monitoring EVC statistics to **cfm** or **vpls** using the **evc-protocol** statement and its options at the **[edit protocols oam ethernet evcs]** hierarchy level.

You can set the number of remote UNIs in the EVC using the **remote-uni-count** *number* statement at the **[edit protocols oam ethernet evcs evcs-protocol]** hierarchy level. The **remote-uni-count** defaults to 1. Configuring a value greater than 1 makes the EVC multipoint-to-multipoint. If you enter a value greater than the actual number of endpoints, the EVC status will display as partially active even if all endpoints are up. If you enter a **remote-uni-count** less than the actual number of endpoints, the status will display as active, even if all endpoints are not up.

You can configure an EVC by including the **evcs** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
evcs evc-id {
    evc-protocol (cfm (management-domain name management-association name ) | vpls
        (routing-instance name)) {
```

```
      remote-uni-count <number>; # Optional, defaults to 1
      multipoint-to-multipoint;
        # Optional, defaults to point-to-point if remote-uni-count is 1
    }
  }
```

Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC

To configure E-LMI, include the **lmi** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
lmi {
  polling-verification-timer value;
    # Polling verification timer (T392), defaults to 15 seconds
  status-counter count; # Status counter (N393), defaults to 4
  interface name {
    evc evc-id {
      default-evc;
      vlan-list [ vlan-ids ];
    }
    evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
    polling-verification-time value; # Optional, defaults to global value
    status-counter count; # Optional, defaults to global value
    uni-id value; # Optional, defaults to interface-name
  }
}
```

You can set the status counter to count consecutive errors using the **status-counter** *count* statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The status counter is used to determine if E-LMI is operational or not. The default value is 4.

You can set the **polling-verification-timer** *value* statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The default value is 15 seconds.

You can enable an interface and set its options for use with E-LMI using the **interface** *name* statement at the **[edit protocols oam ethernet lmi]** hierarchy level. Only **ge**, **xe**, and **ae** interfaces are supported. You can use the interface **uni-id** option to specify a name for the UNI. If **uni-id** is not configured, it defaults to the name variable of **interface** *name*.

You can specify the CE-VLAN ID/EVC map type using the **evc-map-type type** interface option. The options are **all-to-one-bundling**, **bundling**, or **service-multiplexing**. Service multiplexing is with no bundling. The default type is **all-to-one-bundling**.

To specify the EVC that an interface uses, use the **evc** *evc-id* statement at the **[edit protocols oam ethernet lmi interface** *name***]** hierarchy level. You can specify an interface as the default EVC interface using the **default-evc** statement at the **[edit protocols oam ethernet lmi interface** *name* **evc** *evc-id***]** hierarchy level. All VIDs that are not mapped to any other EVCs are mapped to this EVC. Only one EVC can be configured as the default.

You can map a list of VLANs to an EVC using the **vlan-list** *vlan-id-list* statement at the **[edit protocols oam ethernet lmi interface** *name* **evc** *evc-id***]** hierarchy level.

Example E-LMI Configuration

Figure 4 on page 31 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 4: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM

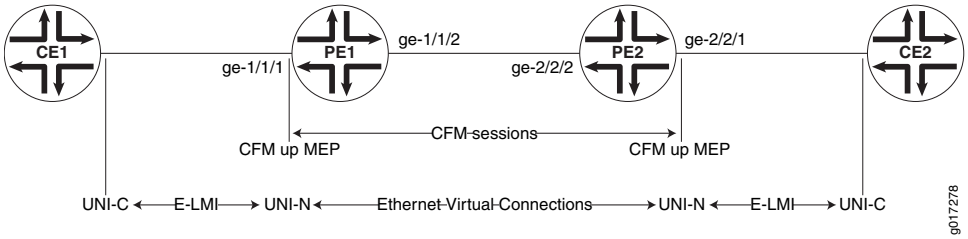

Configuring PE1

```
[edit]
interfaces {
   ge-1/1/1 {
      unit 0 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 1-2048;
         }
      }
      unit 1 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 2049-4096;
         }
      }
   }
   ge-1/1/2 {
      unit 0 {
         vlan-id 100;
         family bridge {
            interface-mode trunk;
            inner-vlan-id-list 1-2048;
         }
      }
      unit 1 {
         vlan-id 200;
         family bridge {
            interface-mode trunk;
            inner-vlan-id-list 2049-4096;
         }
      }
   }
}
protocols {
```

```
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md {
        level 0;
        maintenance-association 1 {
          name-format vlan;
          mep 1 {
            direction up;
            interface ge-1/1/1.0 vlan 1;
          }
        }
        maintenance-association 2049 {
          name-format vlan;
          mep 1 {
            direction up;
            interface ge-1/1/1.1 vlan 2049;
          }
        }
      }
    }
    evcs {
      evc1 {
        evc-protocol cfm management-domain md management-association 1;
        remote-uni-count 1;
      }
      evc2 {
        evc-protocol cfm management-domain md management-association 2049;
        remote-uni-count 1;
      }
    }
    lmi {
      interface ge-1/1/1 {
        evc evc1 {
          vlan-list 1-2048;
        }
        evc evc2 {
          vlan-list 2049-4096;
        }
        evc-map-type bundling;
        uni-id uni-ce1;
      }
    }
  }
}
```

Configuring PE2

```
[edit]
interfaces {
  ge-2/2/1 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 1-2048;
```

```
        }
      }
      unit 1 {
        family bridge {
          interface-mode trunk;
          vlan-id-list 2049-4096;
        }
      }
    }
    ge-2/2/2 {
      unit 0 {
        vlan-id 100;
        family bridge {
          interface-mode trunk;
          inner-vlan-id-list 1-2048;
        }
      }
      unit 1 {
        vlan-id 200;
        family bridge {
          interface-mode trunk;
          inner-vlan-id-list 2049-4095;
        }
      }
    }
}
protocols {
    oam {
      ethernet {
        connectivity-fault-management {
          maintenance-domain md {
            level 0;
            maintenance-association 1 {
              name-format vlan;
              mep 1 {
                direction up;
                interface ge-2/2/1.0 vlan 1;
              }
            }
            maintenance-association 2049 {
              name-format vlan;
              mep 1 {
                direction up;
                interface ge-2/2/1.1 vlan 2049;
              }
            }
          }
        }
        evcs {
          evc1 {
            evc-protocol cfm management-domain md management-association 1;
            remote-uni-count 1;
          }
          evc2 {
            evc-protocol cfm management-domain md management-association 2049;
            uni-count 2;
```

```
        }
      }
      lmi {
        interface ge-2/2/1 {
          evc evc1 {
            vlan-list 1-2048;
          }
          evc evc2 {
            vlan-list 2049-4095;
          }
          evc-map-type bundling;
          uni-id uni-ce2;
        }
      }
    }
  }
}
```

Configuring Two UNIs Sharing the Same EVC

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management { ...}
    evcs {
      evc1 {
        evc-protocol cfm management-domain md management-association 1;
        remote-uni-count 1;
      }
    }
    lmi {
      interface ge-2/2/1 {
        evc evc1 {
          vlan-list 0-4095;
        }
        evc-map-type all-to-one-bundling;
        uni-id uni-ce1;
      }
      interface ge-2/3/1 {
        evc evc1 {
          vlan-list 0-4095;
        }
        evc-map-type all-to-one-bundling;
        uni-id uni-ce2;
      }
    }
  }
}
```

Related Documentation

- connectivity-fault-management on page 103
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13

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 5 on page 35 shows the TLV format and indicates if it is required or optional.

Table 5: 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 6 on page 36 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 6: Type Field Values for Various TLVs for CFM PDUs

TLV or Organization	Type Field
End TLV	0
Sender ID TLV	1
Port Status TLV	2
Data TLV	3
Interface Status TLV	4
Reply Ingress TLV	5
Reply Egress TLV	6
LTM Egress Identifier TLV	7
LTR Egress Identifier TLV	8
Reserved for IEEE 802.1	9 to 30
Organization-Specific TLV	31
Defined by ITU-T Y.1731	32 to 63
Reserved for IEEE 802.1	64 to 255

Not every TLV is applicable for all types of CFM PDUs.

- TLVs applicable for continuity check message (CCM):
  - End TLV
  - Sender ID TLV
  - Port Status TLV
  - Interface Status TLV
  - Organization-Specific TLV
- TLVs applicable for loopback message (LBM):

    - End TLV
    - Sender ID TLV
    - Data TLV
    - Organization-Specific TLV
- TLVs applicable for loopback reply (LBR):
    - End TLV
    - Sender ID TLV
    - Data TLV
    - Organization-Specific TLV
- TLVs applicable for linktrace message (LTM):
    - End TLV
    - LTM Egress Identifier TLV
    - Sender ID TLV
    - Organization-Specific TLV
- TLVs applicable for linktrace reply (LTR):
    - End TLV
    - LTR Egress Identifier TLV
    - Reply Ingress TLV
    - Reply Egress TLV
    - Sender ID TLV
    - Organization-Specific TLV

The following TLVs are currently supported in the applicable CFM PDUs:

- End TLV
- Reply Ingress TLV
- Reply Egress TLV
- LTR Egress Identifier TLV
- LTM Egress Identifier TLV
- Data TLV

Support for Additional Optional TLVs

The following additional optional TLVs are supported:

- Port Status TLV

- Interface Status TLV

MX Series routers support configuration of port status TLV and interface status TLV. Configuring the Port Status TLV allows the operator to control the transmission of the Port Status TLV in CFM PDUs.

NOTE: Although Port Status TLV configuration statements are visible in the CLI on M120 and M320 routers, Port Status TLV cannot be configured on these systems. Port Status TLV can be enabled on a MEP interface only if it is a bridge logical interface, which is not possible on these systems.

For configuration information, see the following sections:

- Port Status TLV on page 38
- Interface Status TLV on page 40

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 8 on page 38. The format of this TLV is shown in Table 7 on page 38.

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

Table 7: Port Status TLV Format

Parameter	Octet (Sequence)
Type = 2	1
Length	2–3
Value (See Table 8 on page 38)	4

Table 8: Port Status TLV Values

Mnemonic	Ordinary Data Passing Freely Through the Port	Value
psBlocked	No: **enableRmepDefect** = false	1
psUp	Yes: **enableRmepDefect** = true	2

The MEP variable **enableRmepDefect** is a boolean variable indicating whether frames on the service instance monitored by the maintenance associations if this MEP are enabled to pass through this bridge port by the Spanning Tree Protocol and VLAN topology management. It is set to TRUE if:

- The bridge port is set in a state where the traffic can pass through it.

- The bridge port is running multiple instances of the spanning tree.
- The MEP interface is not associated with a bridging domain.

Configuring Port Status TLV

Junos OS provides configuration support for the Port Status TLV, allowing you to control the transmission of this TLV in CCM PDUs. The Junos OS provides this configuration at the continuity-check level. By default, the CCM does not include the Port Status TLV. To configure the Port Status TLV, use the **port-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check]** hierarchy level.

NOTE: Port Status TLV configuration is not mandated by IEEE 802.1ag. The Junos OS provides it in order to give more flexibility to the operator; however it receives and processes CCMs with a Port Status TLV, regardless of this configuration.

An example of the configuration statements follows:

```
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                maintenance-domain identifier {
                level number;
                    maintenance-association identifier {
                        continuity-check {
                            interval number,
                            loss-threshold number;
                            hold-interval number;
                            port-status-tlv; # Sets Port Status TLV
                        }
                    }
                }
            }
        }
    }
}
```

You cannot enable Port Status TLV transmission in the following two cases:

- If the MEP interface under the maintenance-association is not of type bridge.
- If the MEP is configured on a physical interface.

Displaying the Received Port Status TLV

The Junos OS saves the last received Port Status TLV from a remote MEP. If the received Port Status value does not correspond to one of the standard values listed in Table 8 on page 38, then the **show** command displays it as "unknown." You can display the last saved received Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier***

maintenance-association *identifier* **local-mep** *identifier* **remote-mep** *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none # RX PORT STATUS
    Interface status TLV: none
```

Displaying the Transmitted Port Status TLV

The Junos OS saves the last transmitted Port Status TLV from a local MEP. If the transmission of the Port Status TLV has not been enabled, then the **show** command displays "none." You can display the last saved transmitted Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** *identifier* **maintenance-association** *identifier* **local-mep** *identifier* **remote-mep** *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up # TX PORT STATUS
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

Interface Status TLV

The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the IETF RFC 2863 IF-MIB. The format of this TLV is shown in Table 9 on page 41. The enumerated values are shown in Table 10 on page 41.

Table 9: Interface Status TLV Format

Parameter	Octet (Sequence)
Type = 4	1
Length	2–3
Value (See Table 10 on page 41)	4

Table 10: Interface Status TLV Values

Mnemonic	Interface Status	Value
isUp	up	1
isDown	down	2
isTesting	testing	3
isUnknown	unknown	4
isDormant	dormant	5
isNotPresent	notPresent	6
isLowerLayerDown	lowerLayerDown	7

NOTE: When the operational status of a logical interface changes from the down state (status value of 2) to the lower layer down state (status value of 7) and vice versa, the LinkDown SNMP trap is not generated. For example, if you configure an aggregated Ethernet interface bundle with a VLAN tag and add a physical interface that is in the operationally down state to the bundle, the operational status of the aggregated Ethernet logical interface bundle at that point is lower layer down (7). If you take the MIC associated with the interface offline, the LinkDown trap is not generated when the logical interface shifts from the lower layer down state to the down state.

Similarly, consider another sample scenario in which an physical interface is added to an aggregated Ethernet bundle that has VLAN tagging and the aggregated Ethernet logical interface is disabled. When the logical interface is disabled, the operational status of the logical interface changes to down. If you disable the physical interface that is part of the aggregated Ethernet bundle, the operational status of the aggregated Ethernet logical interface remains down. If you reenable the aggregated Ethernet logical interface, the operational status of it changes from down to lower layer down. The LinkDown SNMP trap is not generated at this point.

Configuring Interface Status TLV

The Junos OS provides configuration support for the Interface Status TLV, thereby allowing operators to control the transmission of this TLV in CCM PDUs through configuration at the continuity-check level.

NOTE: This configuration is not mandated by IEEE 802.1ag; rather it is provided to give more flexibility to the operator. The Junos OS receives and processes CCMs with the Interface Status TLV, regardless of this configuration.

The interface status TLV configuration is shown below:

```
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                maintenance-domain identifier {
                    level number;
                    maintenance-association identifier {
                        continuity-check {
                            interval number;
                            loss-threshold number;
                            hold-interval number;
                            interface-status-tlv; # Sets the interface status TLV
                        }
                    }
                }
            }
        }
    }
}
```

NOTE: The Junos OS supports transmission of only three out of seven possible values for the Interface Status TLV. The supported values are 1, 2, and 7. However, the Junos OS is capable of receiving any value for the Interface Status TLV.

Displaying the Received Interface Status TLV

The Junos OS saves the last received Interface Status TLV from the remote MEP. If the received Interface Status value does not correspond to one of the standard values listed in Table 9 on page 41, then the **show** command displays "unknown."

You can display this last saved Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** *identifier* **maintenance-association** *identifier* **local-mep** *identifier* **remote-mep** *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
```

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none # displays the Interface Status TLV state
```

Displaying the Transmitted Interface Status TLV

The Junos OS saves the last transmitted Interface Status TLV from a local MEP. If the transmission of Interface Status TLV has not been enabled, then the **show** command displays "none."

You can display the last transmitted Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** *identifier* **maintenance-association** *identifier* **local-mep** *identifier* **remote-mep** *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001

Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

MAC Status Defects

The Junos OS provides MAC status defect information, indicating that one or more of the remote MEPs is reporting a failure in its Port Status TLV or Interface Status TLV. It indicates “yes” if either some remote MEP is reporting that its interface is not isUp (for example, at least one remote MEPs interface is unavailable), or if all remote MEPs are reporting a Port Status TLV that contains some value other than psUp (for example, all

remote MEPs Bridge Ports are not forwarding data). There are two **show** commands you can use to view the MAC Status Defects indication.

Use the **mep-database** command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-association ma6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM           : no
    Erroneous CCM received                 : no
    Cross-connect CCM received             : no
    RDI sent by some MEP                   : no
    Some remote MEP's MAC in error state   : yes # MAC Status Defects yes/no
  Statistics:
    CCMs sent                              : 1658
    CCMs received out of sequence          : 0
    LBMs sent                              : 0
    Valid in-order LBRs received           : 0
    Valid out-of-order LBRs received       : 0
    LBRs received with corrupted data      : 0
    LBRs sent                              : 0
    LTMs sent                              : 0
    LTMs received                          : 0
    LTRs sent                              : 0
    LTRs received                          : 0
    Sequence number of next LTM request    : 0
    1DMs sent                              : 0
    Valid 1DMs received                    : 0
    Invalid 1DMs received                  : 0
    DMMs sent                              : 0
    DMRs sent                              : 0
    Valid DMRs received                    : 0
    Invalid DMRs received                  : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      200     00:05:85:73:39:4a       ok    xe-5/0/0.0
```

Use the **interfaces** command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: up, Port status TLV: up
    MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                 : no
      Erroneous CCM received                       : no
      Cross-connect CCM received                   : no
      RDI sent by some MEP                         : no
      Some remote MEP's MAC in error state         : yes # MAC Status Defects
```

```
yes/no
    Statistics:
      CCMs sent                                       : 1328
      CCMs received out of sequence                   : 0
      LBMs sent                                       : 0
      Valid in-order LBRs received                    : 0
      Valid out-of-order LBRs received                : 0
      LBRs received with corrupted data               : 0
      LBRs sent                                       : 0
      LTMs sent                                       : 0
      LTMs received                                   : 0
      LTRs sent                                       : 0
      LTRs received                                   : 0
      Sequence number of next LTM request             : 0
      1DMs sent                                       : 0
      Valid 1DMs received                             : 0
      Invalid 1DMs received                           : 0
      DMMs sent                                       : 0
      DMRs sent                                       : 0
      Valid DMRs received                             : 0
      Invalid DMRs received                           : 0
    Remote MEP count: 1
      Identifier    MAC address        State    Interface
        200     00:05:85:73:39:4a       ok    xe-5/0/0.0
```

Configuring Remote MEP Action Profile Support

Based on values of **interface-status-tlv** and **port-status-tlv** in the received CCM packets, a specific action, such as **interface-down**, can be taken using the **action-profile** options. Multiple action profiles can be configured on the router, but only one action profile can be assigned to a remote MEP.

The action profile can be configured with at least one event to trigger the action; but the action will be triggered if any one of these events occurs. It is not necessary for all of the configured events to occur to trigger **action**.

An action-profile can be applied only at the remote MEP level.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile tlv-action {
  event {
    # If interface status tlv with value specified in the config is received
    interface-status-tlv down|lower-layer-down;
    # If port status tlv with value specified in the config is received
    port-status-tlv blocked;
    # If connectivity is lost to the peer */
    adjacency-loss;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
  default-actions interface-down;
  }
```

```
# domains
  maintenance-domain identifier {
    # maintenance domain level (0-7)
    level number;
    # association
    maintenance-association identifier {
      mep identifier {
        interface ge-x/y/z.w;
        remote-mep identifier {
          # Apply the action-profile for the remote MEP
          action-profile tlv-action;
        }
      }
    }
  }
```

Monitoring a Remote MEP Action Profile

You can use the **show oam ethernet connectivity-fault-management mep-database** command to view the action profile status of a remote MEP, as in the following example:

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none  # last status TLVs transmitted
 by the router
  Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 200, State: ok # displays the remote MEP name and state

    MAC address: 00:05:85:73:96:1f, Type: Configured
    Interface: ge-1/0/8.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: lower-layer-down
    Action profile: juniper  # displays remote MEP's action profile identifier
      Last event: Interface-status-tlv lower-layer-down  # last remote MEP event

                                                          # to trigger action
      Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
                              # action occurrence time
```

Related Documentation

- connectivity-fault-management on page 103
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15

Configuring MAC Flush Message Processing in CET Mode

In carrier Ethernet transport (CET) mode, MX Series routers are used as provider edge (PE) routers, and Nokia Siemens Networks A2200 Carrier Ethernet Switches (referred to as E-domain devices) that run standard-based protocols are used in the access side. On the MX Series routers, VPLS pseudowires are configured dynamically through label distribution protocol (LDP). On the E-domain devices, topology changes are detected through connectivity fault management (CFM) sessions running between the E-domain devices and the MX Series PE routers. The MX Series PE routers can bring the carrier Ethernet interface down if there is CFM connectivity loss. This triggers a local MAC flush as well as a targeted label distribution protocol (T-LDP) MAC flush notification that gets sent towards the remote MX Series PEs to trigger MAC flush on them.

In CET inter-op mode, MX Series routers need to interoperate with the Nokia Siemens Networks Ax100 Carrier Ethernet access devices (referred to as A-domain devices) that run legacy protocols. Nokia Siemens Networks A4100 and A8100 devices act as an intermediate between the MX Series PE routers and A-domain devices. These intermediate devices perform interworking function (IWF) procedures so that operations administration management (OAM) sessions can be run between MX Series routers and A-domain devices. There are no VPLS pseudowires between the MX Series PE routers and the Nokia Siemens Networks A4100 and A8100 intermediate devices, so there is no LDP protocol running between the PE routers to send topology change notifications. In order to communicate topology changes, MX Series routers can trigger a MAC flush and propagate it in the core. MX Series routers can use action profiles based upon the connection protection type length value (TLV) event. The action profile brings down the carrier edge logical interface in MX Series PE routers, which will trigger a local MAC flush and also propagate the topology change to the core using LDP notification.

For VPLS there is no end-to-end connectivity monitored. The access rings are independently monitored by running CFM down multiple end points (MEPs) on the working and protection paths for each of the services between the E-domain devices and the MX Series PE routers, and between the A-domain devices and the MX Series PE routers the IWF hosted by the Nokia Siemens Networks A-4100 devices. When there is a connectivity failure on the working path, the Nokia Siemens Networks Ax200 devices

perform a switchover to the protection path, triggering a topology change notification (in the form of TLVs carried in CCM) to be sent on the active path.

Figure 5: CET inter-op Dual Homed Topology

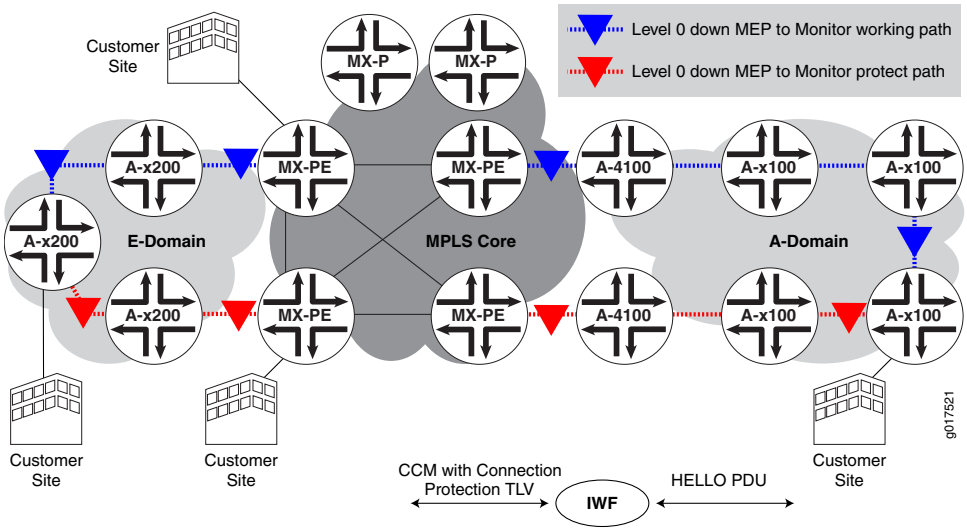

Figure 5 on page 48 describes the dual homed topology on MX Series PE routers connected to the A-domain. When an A-domain device triggers a switchover, it starts switching the service traffic to the new active path. This change is communicated in the HELLO protocol data units (PDUs) sent by that A-domain device on the working and protection paths. When the IWF in A4100 recieves these HELLO PDUs, it converts them to standard CCM messages and also inserts a connection protection TLV. The “Protection-in-use” field of the connection protection TLV is encoded with the currently active path, and is included in the CCM message. CCM messages are received by the MX Series PE routers through the VLAN spoke in A4100. In the above dual homed scenario, one MX Series PE router monitors the working path, and the other MX Series PE router monitors the protection path.

A MAC flush occurs when the CFM session that is monitoring the working path detects that the service traffic has moved to the protection path or when the CFM session that is monitoring the protection path detects that the service traffic has moved to the working path.

Figure 6: CET inter-op Dual Attached Topology

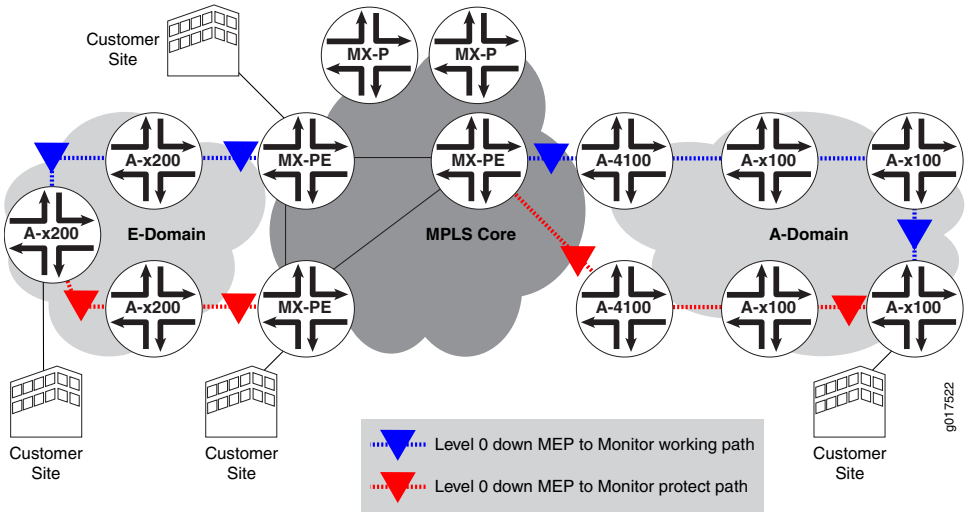

Figure 6 on page 49 describes the dual attached topology on MX Series PE routers connected to the A-domain. The MAC flush mechanism used in this case is also the same as the one used for the A-domain in the dual homed scenario (Figure 1). However in this case both the CFM sessions are hosted by only 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*

Configuring M120 and MX Series Routers for CCC Encapsulated Packets

IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets Overview

Layer 2 virtual private network (L2VPN) is a type of virtual private network service used to transport customer's private Layer 2 traffic (for example, Ethernet, ATM or Frame Relay) over the service provider's shared IP/MPLS infrastructure. The service provider edge (PE) router must have an interface with circuit cross-connect (CCC) encapsulation to switch the customer edge (CE) traffic to the public network.

The IEEE 802.1ag Ethernet Connectivity Fault Management (CFM) is an OAM standard used to perform fault detection, isolation, and verification on virtual bridge LANs. M120 and MX Series routers provide CFM support for bridge/VPLS/routed interfaces and support 802.1ag Ethernet OAM for CCC encapsulated packets.

CFM Features Supported on Layer 2 VPN Circuits

CFM features supported on L2VPN circuits are as follows:

- Creation of up/down MEPs at any level on the CE-facing logical interfaces.
- Creation of MIPs at any level on the CE-facing logical interfaces.
- Support for continuity check, loopback, and linkrace protocol.

- Support for the Y1731 Ethernet Delay measurement protocol.
- Support for action profiles to bring the CE-facing logical interfaces down when loss of connectivity is detected.

Figure 7: Layer 2 VPN Topology

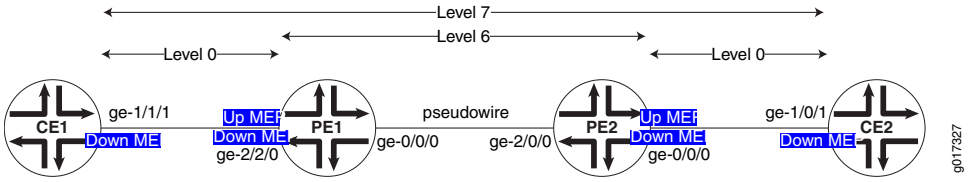

To monitor the L2VPN circuit, a CFM up MEP (Level 6 in Figure 7 on page 51) can be configured on the CE-facing logical interfaces of provider edge routers PE1 and PE2. To monitor the CE-PE attachment circuit, a CFM down MEP can be configured on the customer logical interfaces of CE1-PE1 and CE2-PE2 (Level 0 in Figure 7 on page 51).

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 103
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11

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;
```



```
    }
  }
```

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;
        }
      }
    }
  }
```

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 103**
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Creating the Maintenance Domain on page 11
- Configuring Maintenance Intermediate Points on page 13
- Creating a Maintenance Association on page 15
- Continuity Check Protocol on page 16
- Configuring a Maintenance Endpoint on page 18
- Configuring a Connectivity Fault Management Action Profile on page 23
- Configuring Linktrace Protocol in CFM on page 26
- Configuring Ethernet Local Management Interface on page 27
- Configuring Port Status TLV and Interface Status TLV on page 35
- Configuring MAC Flush Message Processing in CET Mode on page 47
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 50
- *Ethernet Interfaces*

Configuring Unified ISSU for 802.1ag CFM

A unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is automatically enabled for the Connectivity Fault

Management (CFM) protocols and interoperates between local and remote maintenance endpoints (MEPs).

The Junos OS provides support for unified ISSU using the loss threshold type length value (TLV), which is automatically enabled for CFM. TLVs are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and optional information in a protocol data unit (PDU). The loss threshold TLV indicates the loss threshold value of a remote MEP. The loss threshold TLV is transmitted as part of the CFM continuity check messages.

NOTE: **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 3, and you should further observe the additional requirements described in this topic.

Table 11 on page 55 shows the Loss Threshold TLV format.

Table 11: Loss Threshold TLV Format

Parameter	Octet (sequence)	Description
Type=31	1	Required. Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.
Length=12	2	Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field.
OUI	3	Optional. Organization unique identifier (OUI), which is controlled by the IEEE and is typically the first three bytes of a MAC address (Juniper OUI 0x009069).
Subtype	1	Optional. Organizationally defined subtype.
Value	4	Optional. Loss threshold value.
Flag	4	Optional. Bit0 (identifies an ISSU is in progress) Bit1-31 (reserved)

Junos OS provides configuration support for the **convey-loss-threshold** statement, allowing you to control the transmission of the loss threshold TLV in continuity check messages PDUs. The **convey-loss-threshold** statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the **convey-loss-threshold** statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress. The Junos OS provides this configuration at the continuity-check level. By default, continuity check messages do not include the loss threshold TLV.

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

For the remote MEP, the loss threshold TLV is transmitted only during the unified ISSU if the **convey-loss-threshold** statement is not configured. The remote MEP switches back to the default loss threshold if no loss threshold TLV is received or the TLV has a default threshold value of 3.

An example of the ISSU configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
        level number;
          maintenance-association identifier {
            continuity-check {
              convey-loss-threshold;
              interval number;
              loss-threshold number;
              hold-interval number;
            }
          }
        }
      }
    }
  }
}
```

The Junos OS saves the last received loss threshold TLV from the remote MEP. You can display the last saved loss threshold TLV that is received by the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md3 maintenance-association ma5 local-mep 2 remote-mep 1
Maintenance domain name: md3, Format: string, Level: 3
  Maintenance association name: ma3, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
```

```
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
    Loss Threshold TLV:
      Loss Threshold: 3  , Flag:    0x0

 Remote MEP identifier: 1, State: ok
    MAC address: 00:1f:12:b7:ce:79, Type: Learned
    Interface: xe-4/1/1.0
    Last flapped: Never
    Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
    Effective loss threshold: 3 frames
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
    Connection Protection TLV:
      Prefer me: no, Protection in use: no, FRR Flag: no
    Loss Threshold TLV:   #Displays last received value
      Loss Threshold: 3  , Flag:    0x0
```

The Junos OS saves the last transmitted loss threshold TLV from a local MEP. You can display the last transmitted loss threshold TLV and the effective loss (operational) threshold for the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md3 maintenance-association ma5 local-mep 2 remote-mep 1
Maintenance domain name: md3, Format: string, Level: 3
  Maintenance association name: ma3, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
    Loss Threshold TLV:   #Displays last transmitted value
      Loss Threshold: 3  , Flag:    0x0

   Remote MEP identifier: 1, State: ok
    MAC address: 00:1f:12:b7:ce:79, Type: Learned
    Interface: xe-4/1/1.0
    Last flapped: Never
    Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
    Effective loss threshold: 3 frames   #Displays operational threshold
    Remote defect indication: falsePort status TLV: none
    Interface status TLV: none
    Connection Protection TLV:
      Prefer me: no, Protection in use: no, FRR Flag: no
    Loss Threshold TLV:
      Loss Threshold: 3  , Flag:    0x0
```

Related Documentation

- Example: Configuring Ethernet CFM over VPLS on page 64
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3

Configuring Faster Protection Switching

You can apply an action profile to provide faster protection switching for point-to-point network topologies with local switching configured. In a normal state, CCM sessions are configured on the working and protect interfaces. The CCM packets transmitted contain an interface-status TLV with the value up on the working interface and value down on the protect interface. When a link fails on the working interface, the protect interface starts receiving the interface-status TLV as up. With the profile configuration, if the interface-status TLV received on the protect interface is up, the working interface is automatically marked as **interface-down**.

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

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

To configure **interface-down *peer-interface*** as the clear-action, include **interface-down *peer-interface*** at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* clear-action]** hierarchy level.

```
[edit protocols oam]
ethernet {
   connectivity-fault-management {
      action-profile p1 {
         event {
            interface-status-tlv down;
         }
         action {
            interface-down;
         }
         clear-action {
            interface-down peer-interface;
         }
      }
   }
}
```

In this action profile configuration, when the interface-status TLV is received as up, the *peer-interface* is marked as down.

The *peer-interface* is configured in the **protect-maintenance-association** statement. Consider the following example using the **protect-maintenance-association** statement in the configuration:

```
[edit protocols oam]
ethernet {
   connectivity-fault-management {
      action-profile p1 {
         event {
            adjacency-loss;
```

```
        }
        action {
          interface-down;
        }
        clear-action {
          interface-down peer-interface;
        }
      }
      maintenance-domain nsn {
        level 5;
        maintenance-association ma1 {
          protect-maintenance-association ma2;
          continuity-check {
            interval 100ms;
            connection-protection-tlv;
          }
          mep 100 {
            interface ge-1/1/0.0;
            direction down;
            auto-discovery;
          }
        }
        maintenance-association ma2 {
          continuity-check {
            interval 100ms;
            connection-protection-tlv;
        }
        mep 101 {
            interface ge-1/2/0.0;
            direction down;
            auto-discovery;
        }
        remote-mep 100
          action-profile p1;
        }
      }
    }
  }
```

Related Documentation

- connectivity-fault-management on page 103
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- Configuring a Connectivity Fault Management Action Profile on page 23

Configuring Faster Convergence

You can apply an action profile to provide faster convergence for dual-homed multipoint-to-multipoint network topologies. If a multipoint-to-multipoint Ethernet service uses MAC-based forwarding and stale MAC addresses exist in the learning tables, this can result in traffic black holes in the network where incoming traffic is silently discarded, without informing the source that the data did not reach its intended recipient. With the profile configuration, if the interface-status TLV received on the protect interface is up, then the interface-status TLV on the working interface is marked as down and the

PE device for the protect interface propagates a remote MAC-flush message to the PE devices in the virtual private LAN service (VPLS) by using TLDP-MAC-FLUSH. The MAC flush avoids traffic blackholing due to stale mac-db entries.

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

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

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

```
[edit protocols oam]
ethernet {
   connectivity-fault-management {
      action-profile test {
         event {
            interface-status-tlv down;
         }
         action {
            propagate-remote-mac-flush;
         }
         clear-action {
            propagate-remote-mac-flush;
         }
      }
   }
}
```

In this action profile configuration, when the incoming CCM packet contains the interface-status TLV with value down, the **propagate-remote-mac-flush** action is triggered for the action-profile.

Related Documentation

- Configuring MAC Flush Message Processing in CET Mode on page 47
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- **connectivity-fault-management on page 103**
- Configuring a Connectivity Fault Management Action Profile on page 23

Configuring a Primary VLAN ID for Increased Flexibility

You can assign a primary virtual LAN (VLAN) ID in the maintenance association for increased flexibility in the number of tags. When a **vlan-range** or **vlan-id-list** is configured on an interface, the service OAM must run on one of the VLANs. The VLAN assigned for service monitoring is considered the primary VLAN. If a **primary-vid** is not configured, Junos OS assigns the first VLAN from the **vlan-range** or **vlan-id-list**. In earlier releases, Junos OS assigned VLAN 4095.

To configure a primary VLAN ID, you can specify the **primary-vid** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
maintenance domain md3 {
  level 3;
  maintenance-association ma3 {
    primary-vid 2000;
    continuity-check {
      interval 10ms;
      connection-protection-tlv;
    }
    mep 2 {
      interface ge-2/2/0.0;
      direction up;
      auto-discovery;
    }
  }
}
```

Related Documentation

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3
- *connection-protection-tlv*
- Creating a Maintenance Association on page 15
- connectivity-fault-management on page 103

Configuring a Remote Maintenance Association to Accept 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 3
- Creating a Maintenance Association on page 15
- **connectivity-fault-management on page 103**
- *connection-protection-tlv*

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

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 8 on page 63

Figure 8: Topology of CET network

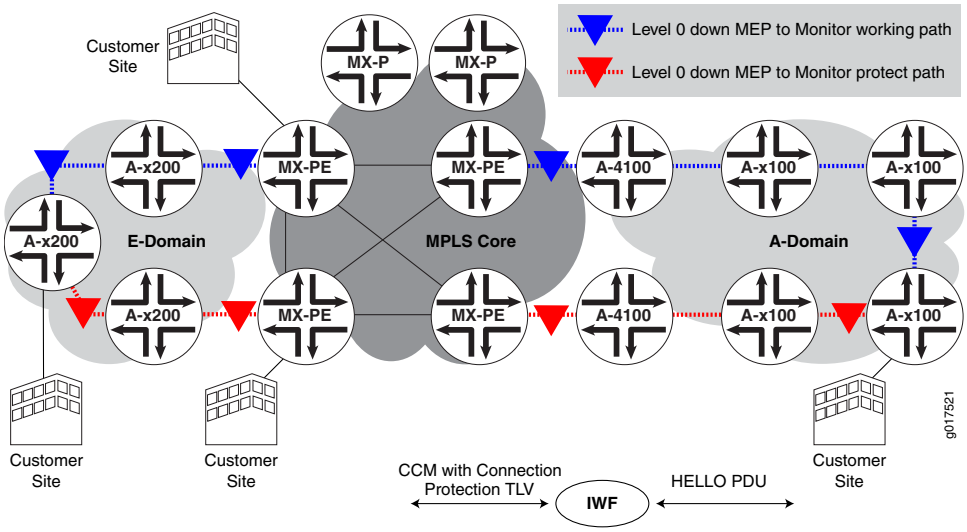

The following definitions describe the meaning of the device abbreviation and terms used in Figure 8 on page 63.

- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.
- E-domain—Nokia Siemens Networks Carrier Ethernet Switches that run standard based protocols and are used in the access side.
- A-domain—Nokia Siemens Networks Carrier Ethernet Switches that run legacy protocols.

Configuration

Step-by-Step Procedure

To configure an action profile based on the connection protection TLV, preform these tasks:

1. Configure an action profile

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
```

2. If the connection protection TLV is received with a “Protection-in-use” value of SET, then the connection protection TLV should use the protection path

```
connection-protection-tlv <using-protection-path>;
```

3. If the connection protection TLV is received with a “Protection-in-use” value of RESET, then the connection protection TLV should use the working path

```
connection-protection-tlv <using-working-path>;
}
```

4. Configure the action profile to bring the interface down

```
action {
  /* Bring the interface down */
```

```
    interface-down;
  }
  }
```

Results Check the results of the configuration

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
    connection-protection-tlv <using-protection-path>;
    connection-protection-tlv <using-working-path>;
  }
  action {
    interface-down;
  }
}
```

Related Documentation

- *connection-protection-tlv*
- Configuring MAC Flush Message Processing in CET Mode on page 47

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 9 on page 65. 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 9: Ethernet OAM with VPLS

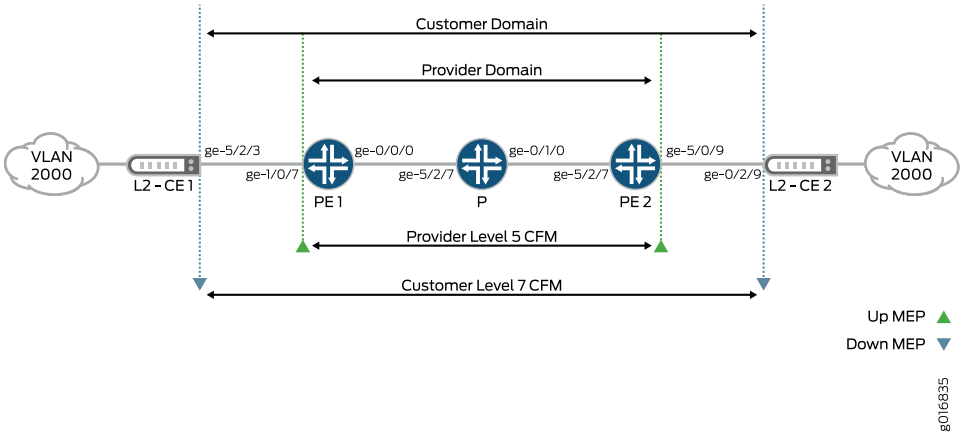

NOTE: The logical interfaces in a VPLS routing instance might have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. Normalization supports automatic mapping of VLANs and performs operations on VLAN tags to achieve the desired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

NOTE:

The following forwarding path considerations must be observed:

- Packet receive path:
  - This is the forwarding path for packets received on the interfaces.
  - 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.
- Packet transmit path:
  - Junos OS uses the router's hardware-based forwarding for CPU-generated packets.
  - For down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - In MX series routers, for up MEPs, the packets must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packets to be forwarded with this flood route.

The following are the configurations of the VPLS and CFM on the service provider routers.

Configuration of PE1 [edit chassis]

```
fpc 5 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}

[edit interfaces]
ge-1/0/7 {
  encapsulation flexible-ethernet-services;
  vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 2000;
  }
}
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.200.1.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.168.231/32 {
        primary;
      }
      address 127.0.0.1/32;
    }
  }
}

[edit routing-instances]
vpls-vlan2000 {
  instance-type vpls;
  vlan-id 2000;
  interface ge-1/0/7.1;
  route-distinguisher 10.255.168.231:2000;
  vrf-target target:1000:1;
  protocols {
    vpls {
      site-range 10;
      site vlan2000-PE1 {
        site-identifier 2;
      }
    }
  }
}

[edit protocols]
rsvp {
```

```
    interface ge-0/0/0.0;
}
mpls {
    label-switched-path PE1-to-PE2 {
        to 10.100.1.1;
    }
    interface ge-0/0/0.0;
}
bgp {
    group PE1-to-PE2 {
        type internal;
        local-address 10.200.1.1;
        family l2vpn {
            signaling;
        }
        local-as 65000;
        neighbor 10.100.1.1;
    }
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface ge-0/0/0.0;
    }
}
oam {
    ethernet {
        connectivity-fault-management {
            maintenance-domain customer-site1 {
                level 5;
                maintenance-association customer-site1 {
                    continuity-check {
                        interval 1s;
                    }
                    mep 100 {
                        interface ge-1/0/7.1;
                        direction up;
                        auto-discovery;
                    }
                }
            }
        }
    }
}
```

Configuration of PE2

```
[edit chassis]
fpc 5 {
    pic 0 {
        tunnel-services {
            bandwidth 1g;
        }
```

```
    }
}

[edit interfaces]
ge-5/0/9 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
        encapsulation vlan-vpls;
        vlan-id 2000;
    }
}
ge-5/2/7 {
    unit 0 {
        family inet {
            address 10.100.1.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.168.230/32 {
                primary;
            }
            address 127.0.0.1/32;
        }
    }
}

[edit routing-instances]
vpls-vlan2000 {
    instance-type vpls;
    vlan-id 2000;
    interface ge-5/0/9.1;
    route-distinguisher 10.255.168.230:2000;
    vrf-target target:1000:1;
    protocols {
        vpls {
            site-range 10;
            site vlan2000-PE2 {
                site-identifier 1;
            }
        }
    }
}

[edit protocols]
rsvp {
    interface ge-5/2/7.0;
}
mpls {
    label-switched-path PE2-to-PE1 {
        to 10.200.1.1;
```

```
    }
    interface ge-5/2/7.0;
  }
  bgp {
    group PE2-to-PE1 {
      type internal;
      local-address 10.100.1.1;
      family l2vpn {
        signaling;
      }
      local-as 65000;
      neighbor 10.200.1.1;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
      interface ge-5/2/7.0;
    }
  }
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain customer-site1 {
          level 5;
          maintenance-association customer-site1 {
            continuity-check {
              interval 1s;
            }
            mep 200 {
              interface ge-5/0/9.1;
              direction up;
              auto-discovery;
            }
          }
        }
      }
    }
  }
```

Configuration of P router

MPLS only, no CFM needed:

```
[edit]
interfaces {
  ge-5/2/7 {
    # Connected to PE1
    unit 0 {
      family inet {
        address 10.200.1.10/24;
      }
      family mpls;
```

```
    }
  }
  ge-0/1/0 {
    # Connected to PE2
    unit 0 {
      family inet {
        address 10.100.1.10/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0{
      family inet {
        address 10.255.168.240/32;
      }
    }
  }
}

[edit]
protocols {
  rsvp {
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
  }
  mpls {
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
      interface ge-0/1/0.0;
      interface ge-5/2/7.0;
    }
  }
}
```

CFM on L2-CE1 Here is the configuration of CFM on L2-E1:

```
[edit interfaces]
ge-5/2/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
```



```
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 800 {
          interface ge-5/2/3.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
```

CFM on L2-CE2 Here is the configuration of CFM L2-CE2:

```
[edit interfaces]
ge-0/2/9 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 700 {
          interface ge-0/2/9.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
```

Related Documentation

- *Ethernet 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*

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

```
    }
}
```

Related Documentation

- Configuring CCM for Better Scalability on page 7
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

[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.

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;
    }
    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);
  }
  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;
  }
  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;
    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;
      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;
  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);
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;
  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 {
        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*
- *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 {
    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;
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;
    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;
                  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*
- *Junos OS Network Interfaces Library for Routing Devices*

[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;
      }
    }
    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
    }
  }
}
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

CHAPTER 4

Statement Summary

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 18

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 23
- default-actions on page 106
- *event (CFM)*
- interface-down on page 113

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 26

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 13
- maintenance-domain on page 119

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;
```

```
            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 3

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 16

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.

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 23

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 18
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 3

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;
```

```
            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*

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 27
- lmi (Ethernet OAM) on page 115

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*

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 16
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 13
- maintenance-domain on page 119

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 23

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 45

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 12
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 26

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.

Options

status-counter *count*—Status counter (N393), defaults to 4.

interface *name*—Polling verification timer (T392), defaults to 15 seconds.

uni-id *uni-name*—(Optional) Defaults to the physical interface name.

status-counter *number*—(Optional) Defaults to a global value.

polling-verification-timer *value*—(Optional) Defaults to a global value.

evc-map-type (all-to-one-bundling | bundling | service-multiplexing)—Specify the Ethernet virtual connection (EVC) map type.

evc *evc-name*—Specify the name of the EVC.

default-evc—Set the specified EVC as the default EVC.

vlan-list *vlan-id-list*—Specify a group of VLANs to assign to the EVC.

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

Related Documentation

- Configuring Ethernet Local Management Interface on page 27
- evcs on page 110

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 16
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 20

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 15
- Configuring a Maintenance Endpoint on page 18
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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.

Related Documentation

- Creating the Maintenance Domain on page 11
- Configuring a Maintenance Endpoint on page 18
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 18

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 11
- maintenance-domain on page 119

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 15
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 (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;
                    }
                  }
                }
              }
            }
          }
          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.

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*
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 26

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 52
- policer (CFM Global) on page 127
- policer (CFM Session) on page 128

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 52
- policer (CFM Session) on page 128

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 52
- policer (CFM Global) on page 127

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 23
- Configuring Remote MEP Action Profile Support on page 45

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 18

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 18

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 18

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 18

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*

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 15
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 72

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 13

PART 3

Administration

CHAPTER 5

Monitoring Commands

- clear oam ethernet connectivity-fault-management continuity-measurement
- clear oam ethernet connectivity-fault-management loss-statistics
- clear oam ethernet connectivity-fault-management policer
- show interfaces (Aggregated Ethernet)
- show interfaces demux0 (Demux Interfaces)
- show interfaces (Fast Ethernet)
- show interfaces (10-Gigabit Ethernet)
- show interfaces interface-set (Ethernet Interface Set)
- 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 mep-database
- show oam ethernet connectivity-fault-management mep-statistics
- show oam ethernet connectivity-fault-management path-database
- show oam ethernet evc
- show oam ethernet link-fault-management
- show oam ethernet lmi
- show oam ethernet lmi statistics

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*
- *Ethernet Interfaces*

List of Sample Output clear oam ethernet connectivity-fault-management continuity-measurement on page 138

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management continuity-measurement

```
user@host> clear oam ethernet connectivity-fault-management continuity-measurement
maintenance-domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
Continuity measurement restarted.
```

clear oam ethernet connectivity-fault-management 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*
- *Displaying ETH-LM Statistics*
- *Managing ETH-LM Statistics*

List of Sample Output clear oam ethernet connectivity-fault-management loss-statistics on page 139

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management loss-statistics

```
user@host> clear oam ethernet connectivity-fault-management loss-statistics
Cleared loss measurements statistics of all CFM sessions
```

clear oam ethernet connectivity-fault-management policer

Syntax clear oam ethernet connectivity-fault-management policer maintenance-domain *md-name* maintenance-association *ma-name*

Release Information Command introduced in Junos OS Release 10.0.

Description On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, M120, MX Series, T320, and T640 routers, clear connectivity-fault-management policer statistics.

Options The following options are supported:

maintenance-domain md-name—Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are cleared for all maintenance associations for all maintenance domains.

maintenance-association ma-name —Name of an existing CFM maintenance association. If this option is not specified, policer statistics are cleared for all maintenance associations for given maintenance domain. This option cannot be specified without specifying maintenance-domain name.

Required Privilege Level view

Related Documentation

- *show oam ethernet connectivity-fault-management policer*

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
```

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 146
show interfaces brief (Aggregated Ethernet) on page 147
show interfaces detail (Aggregated Ethernet) on page 147
show interfaces extensive (Aggregated Ethernet) on page 148
show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 149

Output Fields

Table 12 on page 141 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 12: 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

Table 12: 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**

Table 12: 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**

Table 12: 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

Table 12: 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.	

Table 12: 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
```

```
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
```

```
  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
```

```
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
```

```
    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 158
show interfaces (PPPoE over Aggregated Ethernet) on page 159
show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links) on page 159
show interfaces demux0 (ACI Interface Set Configured) on page 160

Output Fields

Table 13 on page 152 lists the output fields for the **show interfaces** (demux interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 13: 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

Table 13: 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**

Table 13: 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

Table 13: 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 13: Demux show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
protocol-family	Protocol family configured on the logical interface.	**brief**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set. • **IPv6 transit statistics**—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. ◦ **Input bytes**—Number of bytes received on the interface. ◦ **Output bytes**—Number of bytes transmitted on the interface. ◦ **Input packets**—Number of packets received on the interface. ◦ **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	Number of transit bytes and packets received and transmitted on the local interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 Transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. NOTE: The packet and byte counts in these fields include traffic that is dropped and does not leave the router. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input packets	Number of packets received on the interface.	none

Table 13: 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

Table 13: 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
```

```
     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
```

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 174
show interfaces brief (Fast Ethernet) on page 175
show interfaces detail (Fast Ethernet) on page 175
show interfaces extensive (Fast Ethernet) on page 175

Output Fields
Table 14 on page 161 lists the output fields for the **show interfaces Fast Ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 14: show interfaces Fast Ethernet Output Fields

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: 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**

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: 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**

Table 14: 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**

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

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

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 14: 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
```

```
    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
```

```
  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
```

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 192
show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 195
show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 197
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 199
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 199
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 200

Output Fields

See Table 15 on page 178 for the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Table 15: show interfaces Gigabit Ethernet Output Fields

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

Table 15: 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 15: 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 15 on page 178.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 15: 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**

Table 15: 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**

Table 15: 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**—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**

Table 15: 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 16 on page 192 • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

Table 15: 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 15: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 15: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 15: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

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

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

Table 15: 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 15: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 15: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 15: 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 16 on page 192 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 16 on page 192, 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 16: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

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

Sample Output

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

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

```
None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 1024
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:14: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
```

```
    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:
```

```
    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:
```

```
    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
```

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

...

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

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

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

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

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



```
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73: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__
```

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 203
show interfaces interface-set detail on page 203
show interfaces interface-set (ACI Interface Set) on page 204

Output Fields Table 17 on page 202 describes the information for the **show interfaces interface-set** command.

Table 17: 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 17: 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
```

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 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 213
- show oam ethernet connectivity-fault-management mep-database on page 226
- show oam ethernet connectivity-fault-management mep-statistics on page 237

List of Sample Output

show oam ethernet connectivity-fault-management delay-statistics on page 207
show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 207

Output Fields

Table 18 on page 206 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 18: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the **remote-mep** option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.
Two-way delay (usec)	For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.

Sample Output

show oam ethernet connectivity-fault-management delay-statistics

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec

  Remote MEP identifier: 102
  Remote MAC address: 00:04:55:63:39:5a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        29             58
      2        23             59
      3        27             56
      4        29             62
      5        33             68
    Average one-way delay          : 28 usec
    Average one-way delay variation: 3 usec
    Best case one-way delay        : 23 usec
    Worst case one-way delay       : 33 usec
    Average two-way delay          : 60 usec
    Average two-way delay variation: 3 usec
    Best case two-way delay        : 56 usec
    Worst case two-way delay       : 68 usec
```

show oam ethernet connectivity-fault-management delay-statistics remote-mep

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
```

```
  2        273            550
  3        287            571
  4        299            610
  5        313            650
Average one-way delay          : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay        : 259 usec
Worst case one-way delay       : 313 usec
Average two-way delay          : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay        : 519 usec
Worst case two-way delay       : 650 usec
```


show oam ethernet connectivity-fault-management forwarding-state

Syntax **show oam ethernet connectivity-fault-management forwarding-state interface** ***interface-name*** **| instance** ***instance-name*** **<brief | detail | extensive>**

Release Information Command introduced in Junos OS Release 8.4.

Description On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.

Options

interface ***interface-name***—Display forwarding state information for the specified Ethernet interface only.

instance ***instance-name***—Display forwarding state information for the specified forwarding instance only.

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

Required Privilege Level view

List of Sample Output

show oam ethernet connectivity-fault-management forwarding-state instance on page 210
show oam ethernet connectivity-fault-management forwarding-state interface on page 210
show oam ethernet connectivity-fault-management forwarding-state interface detail on page 211
show oam ethernet connectivity-fault-management forwarding-state interfaceinterface-name on page 212

Output Fields Table 19 on page 209 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 19: 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**

Table 19: 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
```

```
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
Interface name: ge-3/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7        down           Receive           none
```

show oam ethernet connectivity-fault-management interfaces

Syntax

show oam ethernet connectivity-fault-management interfaces
<*ethernet-interface-name*>
<**level** *md-level*>
<**brief | detail | extensive**>

Release Information

Command introduced in Junos OS Release 8.4.
Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.
Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for ACX Series and MX Series routers.

Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.

In addition, for Ethernet interfaces on MX Series routers , also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts when **detail** or **extensive** mode is specified.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

brief | detail | extensive—(Optional) Specified level of output.

ethernet-interface-name—(Optional) CFM information only for CFM entities attached to the specified Ethernet interface.

level ***md-level***—(Optional) CFM information for CFM identities enclosed within a maintenance domain of the specified level.

Required Privilege Level

view

Related Documentation

- *clear oam ethernet connectivity-fault-management statistics*
- show oam ethernet connectivity-fault-management delay-statistics on page 205
- show oam ethernet connectivity-fault-management mep-database on page 226
- show oam ethernet connectivity-fault-management mep-statistics on page 237

List of Sample Output

show oam ethernet connectivity-fault-management interfaces (trunk ports) on page 222

Output Fields Table 20 on page 214 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 20: 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**

Table 20: 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**

Table 20: 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**

Table 20: 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**

Table 20: 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
```

```
    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@hostshow oam ethernet connectivity-fault-management interfaces detail

Interface name: xe-6/2/0.0 , Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
 Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: none, Port status TLV: none
    Connection Protection TLV: yes
      Prefer me: no, Protection in use: no, FRR Flag: no
    MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                 : no
      Erroneous CCM received                       : no
      Cross-connect CCM received                   : no
      RDI sent by some MEP                         : no
      Some remote MEP's MAC in error state         : no
    Statistics:
      CCMs sent                                    : 225
      CCMs received out of sequence                : 0
      LBMs sent                                    : 0
      Valid in-order LBRs received                 : 0
      Valid out-of-order LBRs received             : 0
      LBRs received with corrupted data            : 0
      LBRs sent                                    : 0
      LTMs sent                                    : 0
      LTMs received                                : 0
      LTRs sent                                    : 0
      LTRs received                                : 0
      Sequence number of next LTM request          : 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 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
```

```
2001     00:90:69:0b:7f:71       ok    ge-5/2/9.0
4001     00:90:69:0b:09:c5       ok    ge-5/2/9.0
```

show oam ethernet connectivity-fault-management interfaces level

```
user@host> show oam ethernet connectivity-fault-management interfaces level 7
Interface       Link       Status           Level       MEP          Neighbors
                                                        Identifier
ge-3/0/0.0      Up         Active           7           201          0
xe-0/0/0.0      Up         Active           7           203          1
```

show oam ethernet connectivity-fault-management interfaces (trunk ports)

```
user@host> show oam ethernet connectivity-fault-management interfaces

Interface                   Link      Status       Level  MEP          Neighbors
                                                          Identifier
ge-4/0/1.0, vlan 100        Up        Active       5      100          0
ge-10/3/10.4091, vlan 4091  Down      Inactive     4      400          0
ge-4/0/0.0                  Up        Active       6      200          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0

Interface                   Link      Status       Level  MEP          Neighbors
                                                          Identifier
ge-4/0/0.0                  Up        Active       6      200          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vlan 100

Interface                   Link      Status       Level  MEP          Neighbors
                                                          Identifier
ge-4/0/1.0, vlan 100        Up        Active       5      100          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091
vlan 4091

Interface                   Link      Status       Level  MEP          Neighbors
                                                          Identifier
ge-10/3/10.4091, vlan 4091  Down      Inactive     4      400          0
```

show oam ethernet connectivity-fault-management linktrace path-database

Syntax **show oam ethernet connectivity-fault-management linktrace path-database mac-address** *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 224
show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 224

Output Fields Table 21 on page 223 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 21: 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.

Table 21: 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
```

```
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 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 205
- show oam ethernet connectivity-fault-management interfaces on page 213
- show oam ethernet connectivity-fault-management mep-statistics on page 237

List of Sample Output

show oam ethernet connectivity-fault-management mep-database on page 231
show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM) on page 232
show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 233

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 233
show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured) on page 233
show oam ethernet connectivity-fault-management mep-database on page 234
show oam ethernet connectivity-fault-management mep-database (enhanced continuity measurement) on page 235

Output Fields Table 22 on page 227 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 22: 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**.

Table 22: 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**.

Table 22: 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.

Table 22: 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.

Table 22: 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:
```

```
    CCMs sent                                      : 1476
    CCMs received out of sequence                  : 0
    LBMs sent                                      : 85
    Valid in-order LBRs received                   : 78
    Valid out-of-order LBRs received               : 0
    LBRs received with corrupted data              : 0
    LBRs sent                                      : 0
    LTMs sent                                      : 1
    LTMs received                                  : 0
    LTRs sent                                      : 0
    LTRs received                                  : 1
    Sequence number of next LTM request            : 1
    1DMs sent                                      : 0
    Valid 1DMs received                            : 0
    Invalid 1DMs received                          : 0
    DMMs sent                                      : 0
    DMRs sent                                      : 0
    Valid DMRs received                            : 0
    Invalid DMRs received                          : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      100     00:19:e2:b2:81:4b       ok  vt-0/1/10.1049088
```

show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-domain ma6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                   : no
    Erroneous CCM received                         : no
    Cross-connect CCM received                     : no
    RDI sent by some MEP                           : no
  Statistics:
    CCMs sent                                      : 1590
    CCMs received out of sequence                  : 0
    LBMs sent                                      : 0
    Valid in-order LBRs received                   : 0
    Valid out-of-order LBRs received               : 0
    LBRs received with corrupted data              : 0
    LBRs sent                                      : 0
    LTMs sent                                      : 0
    LTMs received                                  : 0
    LTRs sent                                      : 0
    LTRs received                                  : 0
    Sequence number of next LTM request            : 0
    1DMs sent                                      : 10
    Valid 1DMs received                            : 0
    Invalid 1DMs received                          : 0
    DMMs sent                                      : 0
    DMRs sent                                      : 0
    Valid DMRs received                            : 0
    Invalid DMRs received                          : 0
  Remote MEP count: 1
```

```
  Identifier    MAC address        State    Interface
    201     00:90:69:0a:43:94       ok    ge-0/2/5.0
```

show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200
remote-mep 100
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 100, State: ok
    MAC address: 00:19:e2:b2:81:4b, Type: Learned
    Interface: vt-0/1/10.1049088
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 200, State: ok
    MAC address: 00:05:85:73:96:1f, Type: Configured
    Interface: ge-1/0/8.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: lower-layer-down
    Action profile: juniper
      Last event: Interface-status-tlv lower-layer-down
      Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
```

show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured)

```
user@host>show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5
```

If connection-protection is not enabled on down MEPs, but connection-protection TLV is used, MX always sets the protection-in-use flag in connection-protection tlv, while CCMs are sent out. During reversion, this is an indicator to the receiver that protect-path is in use, otherwise the peer (receiver) assumes working is active and reversion does not work as expected. Setting this bit does not affect protection-switching/traffic-loss.

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-6/2/0.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                 : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                         : no
    Some remote MEP's MAC in error state         : no
  Statistics:
    CCMs sent                                    : 251
    CCMs received out of sequence                : 0
    LBMs sent                                    : 0
    Valid in-order LBRs received                 : 0
    Valid out-of-order LBRs received             : 0
    LBRs received with corrupted data            : 0
    LBRs sent                                    : 0
    LTMs sent                                    : 0
    LTMs received                                : 0
    LTRs sent                                    : 0
    LTRs received                                : 0
    Sequence number of next LTM request          : 0
    1DMs sent                                    : 0
    Valid 1DMs received                          : 0
    Invalid 1DMs received                        : 0
    Out of sync 1DMs received                    : 0
    DMMs sent                                    : 0
    Valid DMMs received                          : 0
    Invalid DMMs received                        : 0
    DMRs sent                                    : 0
    Valid DMRs received                          : 0
    Invalid DMRs received                        : 0
    LMMs sent                                    : 0
    Valid LMMs received                          : 0
    Invalid LMMs received                        : 0
    LMRs sent                                    : 0
    Valid LMRs received                          : 0
    Invalid LMRs received                        : 0
  Remote MEP count: 1
   Identifier    MAC address        State    Interface
        2     00:90:69:7f:e4:30
```

show oam ethernet connectivity-fault-management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:14:f6:b6:01:fe
Auto-discovery: enabled, Priority: 0
Interface name: ge-1/0/0.0, Interface status: Active, Link status: Up

Defects:
Remote MEP not receiving CCM         : no
```

```
Erroneous CCM received               : no
Cross-connect CCM received           : no
RDI sent by some MEP                 : no

Statistics:
CCMs sent                            : 328703
CCMs received out of sequence        : 0
LBMs sent                            : 85
Valid in-order LBRs received         : 78
Valid out-of-order LBRs received     : 0
LBRs received with corrupted data    : 0
LBRs sent                            : 0
LTMs sent                            : 0
LTMs received                        : 0
LTRs sent                            : 0
LTRs received                        : 0
Sequence number of next LTM request  : 0
1DMs sent                            : 10
Valid 1DMs received                  : 10
Invalid 1DMs received                : 0
DMMs sent                            : 20
DMRs sent                            : 0
Valid DMRs received                  : 10
Invalid DMRs received                : 0
LMM sent                             : 10
Valid LMM received                   : 20
Invalid LMM received                 : 0
LMR sent                             : 20
Valid LMR received                   : 10
Invalid LMR received                 : 0
SLM sent                                     : 10
Valid SLM received                           : 20
Invalid SLM received                         : 0
SLR sent                                     : 20
Valid SLR received                           : 10
Invalid SLR received                         : 0

Remote MEP count                     : 1

Identifier     MAC address         State     Interface
    2       00:12:1e:fb:ea:7d       ok        ge-1/0/0.0
```

show oam ethernet connectivity-fault-management mep-database (enhanced continuity measurement)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address  : 00:19:e2:b0:74:00, Type: Learned
Interface    : ge-2/0/0.0
Last flapped  : Never
+ Continuity  : 91%, Admin-enable duration: 2100sec, Oper-down duration: 100sec
Remote defect indication: false
```

```
Port status TLV: none
Interface status TLV: none
```

show oam ethernet connectivity-fault-management mep-statistics

Syntax

show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain *md-name*
maintenance-association *ma-name*
<**mep** *mep-id*>
<**remote-mep** *remote-mep-id*>
<**count** *entry-count*>

Release Information

Command introduced in Junos OS Release 9.5.
Command introduced in Junos OS Release 11.4 for EX Series switches.
Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for MX Series routers.

Description

On MX Series and ACX Series routers and EX Series switches with Ethernet interfaces, display ETH-DM statistics and ETH-DM frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

Options

maintenance-domain *md-name*—Name of an existing CFM maintenance domain.

maintenance-association *ma-name*—Name of an existing CFM maintenance association.

mep *mep-id*—(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.

remote-mep *remote-mep-id*—(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.

count *entry-count*—(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.

Required Privilege Level

view

Related Documentation

- *clear oam ethernet connectivity-fault-management statistics*
- show oam ethernet connectivity-fault-management delay-statistics on page 205
- show oam ethernet connectivity-fault-management interfaces on page 213
- show oam ethernet connectivity-fault-management mep-database on page 226

List of Sample Output

show oam ethernet connectivity-fault-management mep-statistics (CIR counters only) on page 239
show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled) on page 241
show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only) on page 242

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled) on page 243
show oam ethernet connectivity-fault-management mep-statistics on page 245
show oam ethernet connectivity-fault-management mep-statistics remote-mep on page 246

Output Fields Table 23 on page 238 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 23: 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.

Table 23: 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
```

```
LBRs sent                           : 0
LTMs sent                           : 0
LTMs received                       : 0
LTRs sent                           : 0
LTRs received                       : 0
Sequence number of next LTM request : 0
1DMs sent                           : 5
Valid 1DMs received                 : 0
Invalid 1DMs received               : 0
DMMs sent                           : 5
DMRs sent                           : 0
Valid DMRs received                 : 5
Invalid DMRs received               : 0
LMM sent                            : 5
Valid LMM received                  : 5
Invalid LMM received                : 0
LMR sent                            : 0
Valid LMR received                  : 5
Invalid LMR received                : 0
Remote MEP identifier               : 101
Remote MAC address                  : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
  4             299                   610
  5             313                   650

Average one-way delay                : 286 usec
Average one-way delay variation      : 62 usec
Best case one-way delay              : 259 usec
Average two-way delay                : 580 usec
Average two-way delay variation      : 26 usec
Best case two-way delay              : 519 usec
Worst case two-way delay             : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)       (CIR)         (EIR)        (EIR)
  1           9           9
  2           3           5
  3           7           5
  4           9           6
  5           3           6

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)     : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)      : 6.2%
Near-end best case loss (CIR)         : 3
Near-end best case loss ratio (CIR)   : 3%
Near-end worst case loss (CIR)        : 9
Near-end worst case loss ratio (CIR)  : 9%
Far-end best case loss (CIR)          : 5
Far-end best case loss ratio (CIR)    : 5%
Far-end worst case loss (CIR)         : 9
Far-end worst case loss ratio (CIR)   : 9%
```

show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count                      : 1
CCMs sent                             : 6550
CCMs received out of sequence         : 0
LBMs sent                             : 0
Valid in-order LBRs received          : 0
Valid out-of-order LBRs received      : 0
LBRs received with corrupted data     : 0
LBRs sent                             : 0
LTMs sent                             : 0
LTMs received                         : 0
LTRs sent                             : 0
LTRs received                         : 0
Sequence number of next LTM request   : 0
1DMs sent                             : 5
Valid 1DMs received                   : 0
Invalid 1DMs received                 : 0
DMMs sent                             : 5
DMRs sent                             : 0
Valid DMRs received                   : 5
Invalid DMRs received                 : 0
LMM sent                              : 5
alid LMM received                     : 5
Invalid LMM received                  : 0
LMR sent                              : 0
Valid LMR received                    : 5
Invalid LMR received                  : 0
Remote MEP identifier                 : 101
Remote MAC address                    : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
              (usec)                (usec)
  1            259                   519
  2            273                  550
  3            287                  571
  4            299                  610
  5            313                  650

Average one-way delay                   : 286 usec
Average one-way delay variation         : 62 usec
Best case one-way delay                 : 259 usec
Average two-way delay                   : 580 usec
Average two-way delay variation         : 26 usec
Best case two-way delay                 : 519 usec
Worst case two-way delay                : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
           (CIR)         (CIR)        (EIR)        (EIR)
  1          9             9            2            4
  2          3             5            4            6
  3          7             5            0            2
  4          9             6            8            2
  5          3             6            6            4
```

```
Average near-end loss (CIR)               : 6.2
Average near-end loss ratio (CIR)         : 6.2%
Average far-end loss (CIR)                : 6.2
Average far-end loss ratio (CIR)          : 6.2%
Near-end best case loss (CIR)             : 3
Near-end best case loss ratio (CIR)       : 3%
Near-end worst case loss (CIR)            : 9
Near-end worst case loss ratio (CIR)      : 9%
Far-end best case loss (CIR)              : 5
Far-end best case loss ratio (CIR)        : 5%
Far-end worst case loss (CIR)             : 9
Far-end worst case loss ratio (CIR)       : 9%
Average near-end loss (EIR)               : 4
Average near-end loss ratio (EIR)         : 4%
Average far-end loss (EIR)                : 3.4
Average far-end loss ratio (EIR)          : 3.4%
Near-end best case loss (EIR)             : 0
Near-end best case loss ratio (EIR)       : 0%
Near-end worst case loss (EIR)            : 8
Near-end worst case loss ratio (EIR)      : 8%
Far-end best case loss (EIR)              : 2
Far-end best case loss ratio (EIR)        : 2%
Far-end worst case loss (EIR)             : 6
Far-end worst case loss ratio (EIR)       : 6%
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                               : 7762
CCMs received out of sequence           : 0
LBMs sent                               : 0
Valid in-order LBRs received            : 0
Valid out-of-order LBRs received        : 0
LBRs received with corrupted data       : 0
LBRs sent                               : 0
LTMs sent                               : 0
LTMs received                           : 0
LTRs sent                               : 0
LTRs received                           : 0
Sequence number of next LTM request     : 0
1DMs sent                               : 5
Valid 1DMs received                     : 0
Invalid 1DMs received                   : 0
DMMs sent                               : 5
DMRs sent                               : 0
Valid DMRs received                     : 5
Invalid DMRs received                   : 0
LMM sent                                : 5
Valid LMM received                      : 5
Invalid LMM received                    : 0
LMR sent                                : 0
Valid LMR received                      : 5
Invalid LMR received                    : 0
Remote MEP identifier                   : 101
Remote MAC address                      : 00:05:85:73:39:4a
```

```
Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
  4             299                   610
  5             313                   650

Average one-way delay                  : 286 usec
Average one-way delay variation        : 62 usec
Best case one-way delay                : 259 usec
Average two-way delay                  : 580 usec
Average two-way delay variation        : 26 usec
Best case two-way delay                : 519 usec
Worst case two-way delay               : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)        (CIR)         (EIR)        (EIR)
  1           9            9
  2           3            5
  3           7            5
  4           9            6
  5           3            6
Average near-end loss (CIR)               : 6.2
Average near-end loss ratio (CIR)         : 6.2%
Average far-end loss (CIR)                : 6.2
Average far-end loss ratio (CIR)          : 6.2%
Near-end best case loss (CIR)             : 3
Near-end best case loss ratio (CIR)       : 3%
Near-end worst case loss (CIR)            : 9
Near-end worst case loss ratio (CIR)      : 9%
Far-end best case loss (CIR)              : 5
Far-end best case loss ratio (CIR)        : 5%
Far-end worst case loss (CIR)             : 9
Far-end worst case loss ratio (CIR)       : 9%
Average near-end loss (EIR)               : 4
Average near-end loss ratio (EIR)         : 4%
Average far-end loss (EIR)                : 3.4
Average far-end loss ratio (EIR)          : 3.4%
Near-end best case loss (EIR)             : 0
Near-end best case loss ratio (EIR)       : 0%
Near-end worst case loss (EIR)            : 8
Near-end worst case loss ratio (EIR)      : 8%
Far-end best case loss (EIR)              : 2
Far-end best case loss ratio (EIR)        : 2%
Far-end worst case loss (EIR)             : 6
Far-end worst case loss ratio (EIR)       : 6%
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled)

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3
remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                           :7762
CCMs received out of sequence       : 0
LBMs sent                           : 0
```

```
Valid in-order LBRs received         : 0
Valid out-of-order LBRs received     : 0
LBRs received with corrupted data    : 0
LBRs sent                            : 0
LTMs sent                            : 0
LTMs received                        : 0
LTRs sent                            : 0
LTRs received                        : 0
Sequence number of next LTM request  : 0
1DMs sent                            : 5
Valid 1DMs received                  : 0
Invalid 1DMs received                : 0
DMMs sent                            : 5
DMRs sent                            : 0
Valid DMRs received                  : 5
Invalid DMRs received                : 0
LMM sent                             : 5
Valid LMM received                   : 5
Invalid LMM received                 : 0
LMR sent                             : 0
Valid LMR received                   : 5
Invalid LMR received                 : 0
Remote MEP identifier                : 101
Remote MAC address                   : 00:05:85:73:39:4a

Delay measurement statistics:
Index       One-way delay         Two-way delay
               (usec)                (usec)
  1             259                   519
  2             273                   550
  3             287                   571
  4             299                   610
  5             313                   650

Average one-way delay                 : 286 usec
Average one-way delay variation       : 62 usec
Best case one-way delay               : 259 usec
Average two-way delay                 : 580 usec
Average two-way delay variation       : 26 usec
Best case two-way delay               : 519 usec
Worst case two-way delay              : 650 usec

Loss measurement statistics:
Index     Near-end     Far-end       Near-end     Far-end
          Frame loss   Frame loss    Frame loss   Frame loss
            (CIR)        (CIR)         (EIR)        (EIR)
  1          10           8             5            12
  2          12           7             6            16
  3          7            5             0            2
  4          9            6             8            2
  5          3            6             6            4

Average near-end loss (CIR)            : 6.2
Average near-end loss ratio (CIR)      : 6.2%
Average far-end loss (CIR)             : 6.2
Average far-end loss ratio (CIR)       : 6.2%
Near-end best case loss (CIR)          : 3
Near-end best case loss ratio (CIR)    : 3%
Near-end worst case loss (CIR)         : 9
Near-end worst case loss ratio (CIR)   : 9%
Far-end best case loss (CIR)           : 5
```

```
Far-end best case loss ratio (CIR)      : 5%
Far-end worst case loss (CIR)           : 9
Far-end worst case loss ratio (CIR)     : 9%
Average near-end loss (EIR)             : 4
Average near-end loss ratio (EIR)       : 4%
Average far-end loss (EIR)              : 3.4
Average far-end loss ratio (EIR)        : 3.4%
Near-end best case loss (EIR)           : 0
Near-end best case loss ratio (EIR)     : 0%
Near-end worst case loss (EIR)          : 8
Near-end worst case loss ratio (EIR)    : 8%
Far-end best case loss (EIR)            : 2
Far-end best case loss ratio (EIR)      : 2%
Far-end worst case loss (EIR)           : 6
Far-end worst case loss ratio (EIR)     : 6%
```

show oam ethernet connectivity-fault-management mep-statistics

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma-1
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 1
  CCMs sent                                     : 6550
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           : 0
  1DMs sent                                     : 5
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
  DMMs sent                                     : 5
  DMRs sent                                     : 0
  Valid DMRs received                           : 5
  Invalid DMRs received                         : 0
  SLM sent                                      : 10
  Valid SLM received                            : 20
  Invalid SLM received                          : 0
  SLR sent                                      : 20
  Valid SLR received                            : 10
  Invalid SLR received                          : 0

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
```

```
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec    Synthetic Loss measurement
statistics:
     SLM packets sent                            : 100
     SLM packets received                        : 0
     SLR packets sent                            : 100
     SLR packets received                        : 0
     Accumulated SLM statistics:
     Local TXFC1 value                           : 100
     Local RXFC1 value                           : 100
     Last Received SLR frame TXFCftc             : 100
     Last Received SLR frame TXFCbtc             : 100
     SLM Frame Loss:
     Frame Loss (far-end)                        : 0 (0.00 %)
     Frame Loss (near-end)                       : 0 (0.00 %)
```

show oam ethernet connectivity-fault-management mep-statistics remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma1 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
  CCMs sent                                     : 7762
  CCMs received out of sequence                 : 0
  LBMs sent                                     : 0
  Valid in-order LBRs received                  : 0
  Valid out-of-order LBRs received              : 0
  LBRs received with corrupted data             : 0
  LBRs sent                                     : 0
  LTMs sent                                     : 0
  LTMs received                                 : 0
  LTRs sent                                     : 0
  LTRs received                                 : 0
  Sequence number of next LTM request           : 0
  1DMs sent                                     : 5
  Valid 1DMs received                           : 0
  Invalid 1DMs received                         : 0
  DMMs sent                                     : 5
  DMRs sent                                     : 0
  Valid DMRs received                           : 5
  Invalid DMRs received                         : 0
  SLM sent                                      : 10
  Valid SLM received                            : 20
  Invalid SLM received                          : 0
  SLR sent                                      : 20
  Valid SLR received                            : 10
  Invalid SLR received                          : 0

  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        259            519
      2        273            550
      3        287            571
      4        299            610
```

```
      5        313            650
    Average one-way delay          : 286 usec
    Average one-way delay variation: 62 usec
    Best case one-way delay        : 259 usec
    Worst case one-way delay       : 313 usec
    Average two-way delay          : 580 usec
    Average two-way delay variation: 26 usec
    Best case two-way delay        : 519 usec
    Worst case two-way delay       : 650 usec    Synthetic Loss measurement
statistics:
     SLM packets sent                           : 100
     SLM packets received                       : 0
     SLR packets sent                           : 100
     SLR packets received                       : 0
     Accumulated SLM statistics:
     Local TXFC1 value                          : 100
     Local RXFC1 value                          : 100
     Last Received SLR frame TXFCftc            : 100
     Last Received SLR frame TXFCbtc            : 100
     SLM Frame Loss:
     Frame Loss (far-end)                       : 0 (0.00 %)
     Frame Loss (near-end)                      : 0 (0.00 %)
```

show oam ethernet connectivity-fault-management 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 249

Output Fields

Table 24 on page 248 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 24: 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.

Table 24: 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 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 25 on page 250 lists the output fields for the **show oam ethernet evc** command. Output fields are listed in the approximate order in which they appear.

Table 25: show oam ethernet evc Output Fields

Field Name	Field Description
EVC identifier	Header for the EVC information showing the EVC name, configuration, and active/inactive status.
UNI count	Number of configured and active UNIs.
Protocol	Protocol configured between the UNIs.
Local UNIs	Heading for the list of local UNIs
UNI Identifier	Name of the UNI.
Interface	Interface type-dpc/pic/port.unit-number.
Status	Status operational or not operational.

Sample Output

show oam ethernet evc

```
user@host> show oam ethernet evc
EVC identifier: evc1, Point-to-Point, Active
  UNI count: Configured(2), Active(2)
  Protocol: cfm, Management domain: md, Management association: ma
  Local UNIs:
    UNI Identifier      Interface     Status
    uni1                ge-1/1/1      Operational
    uni2                ge-1/1/1      Not Operational
```

show oam ethernet 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 255
show oam ethernet link-fault-management detail on page 255

Output Fields Table 26 on page 251 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 26: 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

Table 26: 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**

Table 26: 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**

Table 26: 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		

Table 26: 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
```

```
OAM transmit statistics:
  Information: 186347, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame period error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM current symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM current frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported, Remote in ISSU
```

show oam ethernet lmi

Syntax **show oam ethernet lmi (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, and OAM Ethernet Local Management Interface (LMI) configuration, display the LMI information for the configured interfaces or optionally for a specified interface.

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 27 on page 257 lists the output fields for the **show oam ethernet lmi** command. Output fields are listed in the approximate order in which they appear.

Table 27: 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.
Reference ID	To be provided.
Status	Status active or not active.
CE VLAN IDs	Customer edge VLAN ID numbers.

Sample Output

show oam ethernet lmi interface

```
user@host> show oam ethernet lmi interface ge-1/1/1
Physical interface: ge-1/1/1, Physical link is Up
  UNI identifier: uni-ce1, EVC map type: Bundling
  Polling verification timer: Enabled, E-LMI state: Operational
  Priority/Untagged VLAN ID: 20, Default EVC: evc1
  Associated EVCs:
    EVC         Reference     Status           CE VLAN IDs
    Identifier  ID
    evc1            1        Active (New)        1-2048
    evc2             2        Not Active          2049-4096
```

show oam ethernet lmi statistics

Syntax **show oam ethernet lmi statistics <interface *interface-name*>**

Release Information Command introduced in Junos OS Release 9.5.

Description On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Local Management Interface (LMI) statistics.

Options **interface**—(Optional) Display LMI statistics for a specified interface.

interface-name—(Optional) Display Ethernet LMI information for the specified Ethernet interface only.

Required Privilege Level view

List of Sample Output show oam ethernet lmi statistics on page 259

Output Fields Table 28 on page 259 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 28: show oam ethernet lmi statistics Output Fields

Field Name	Field Description
Physical interface	Name of the interface for the displayed statistics.
Reliability errors	Number of E-LMI reliability errors logged.
Protocol errors	Number of E-LMI protocol errorse.
Status check received	Number of E-LMI status check receive errors.
Status check sent	Number of E-LMI status check sent errors.
Full status received	Number of E-LMI full status receive errors.
Full status sent	Number of E-LMI full status sent errors.
Full status continued received	Number of E-LMI status continued received errors.
Full status continued sent	Number of E-LMI full status continued sent errors.
Asynchronous status sent	Number of E-LMI asynchronous status sent errors.

Sample Output

show oam ethernet lmi statistics

```
user@host> show oam ethernet lmi statistics interface ge-1/1/1
```

```
Physical interface: ge-1/1/1
  Reliability errors                     4  Protocol errors
 0
  Status check received                  0  Status check sent
 0
  Full status received                 694  Full status sent
694
  Full status continued received         0  Full status continued sent
 0
  Asynchronous status sent               0
```

PART 4

Troubleshooting

CHAPTER 6

Ethernet

- traceroute ethernet

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 265

Output Fields Table 29 on page 264 lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 29: 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).

Table 29: 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
```


CHAPTER 7

Interface Diagnostics

Interface Diagnostics

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 the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

There are several types of loopback testing supported by the Junos OS, as follows:

- DCE local—Loops packets back on the local 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 30 on page 268 shows the loopback modes supported on the various interface types.

Table 30: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	*Configuring Ethernet Loopback Capability*
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*
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*

Table 30: Loopback Modes by Interface Type *(continued)*

Interface	Loopback Modes	Usage Guidelines
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]**

Interface Diagnostics

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.

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 O.152 standard)
  pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.152 standard)
  pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
  pseudo-2e20-o153     Pattern is 2^20 - 1 (per O.153 standard)
        ...
```

For specific hierarchy information, see the individual interface types.

NOTE: The 4-port E1 PIC supports only the following algorithms:

```
pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e23-o151     Pattern is 2^23 (per O.151 standard)
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating   Repeating one bits
all-zeros-repeating  Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e7           Pattern is 2^7 - 1
pseudo-2e9-o153      Pattern is 2^9 - 1 (per O.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 O.153 (511 type) standard)
pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 and O.153 (2047 type)
 standards)
pseudo-2e15-o151     Pattern is 2^15 -1 (per O.151 standard)
pseudo-2e20-o151     Pattern is 2^20 -1 (per O.151 standard)
pseudo-2e20-o153     Pattern is 2^20 -1 (per O.153 standard)
pseudo-2e23-o151     Pattern is 2^23 -1 (per O.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 31 on page 272 shows the BERT capabilities for various interface types.

Table 31: 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	N/A	Yes (channel 0–11)	• Single channel at a time • Limited algorithms • No bit count
Channelized STM1	Yes (channel 0–62)	N/A	• Multiple channels • Only one algorithm • No error insert • No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	• Multiple ports and channels • Limited algorithms for T1 • No error insert for T1 • No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see *Channelized IQ and IQE Interfaces Properties*.

Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
disable;
```

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name interface-type*-bert-start** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you wish 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 a 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;
```

```
    }
}
```

PART 5

Index

Index

Symbols

A

B

C

D

E