Junos® OS

Aggregated Ethernet Interfaces 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 Aggregated Ethernet Interfaces 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

Part 4 Troubleshooting

Part 5 Index

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
- PTX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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


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

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

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

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

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

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

Documentation Conventions

Table 1 on page xv 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 xvi 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

Aggregated Ethernet Interfaces

Aggregated Ethernet Interfaces Overview

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

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

Platform Support for Aggregated Ethernet Interfaces

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

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

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

NOTE:

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

TIP:

Going forward:

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

Table 3 on page 4 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 3: Platform Support Matrix for Mixed Aggregated Ethernet Bundles

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
10-Gigabit Ethernet LAN and WAN (WAN rate: OC192)	T640, T1600, T4000, and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (PF-24XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
		• Enhanced Scaling FPC3 (T640-FPC3-ES)	• 10-Gigabit Ethernet PIC with XENPAK (PC-1XGE-XENPAK)
		• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)

Table 3: Platform Support Matrix for Mixed Aggregated Ethernet Bundles *(continued)*

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
40-Gigabit Ethernet, 100-Gigabit Ethernet	T4000 and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 100-Gigabit Ethernet PIC with CFP (PF-1CGE-CFP)
	T640, T1600, T4000, and TX Matrix Plus routers	• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4) NOTE: This PIC is available packaged only in an assembly with the T1600-FPC4-ES FPC. • 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP)

The following routers support a maximum of 16 physical interfaces per single aggregated Ethernet bundle:

- M120
- M320
- All MX Series 3D Universal Edge Routers
- All T Series routers

All other routers support a maximum of eight physical interfaces per aggregated Ethernet bundle.

On M Series and T Series routers, you can create a maximum of 1024 logical interfaces on an aggregated Ethernet interface.

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

Configuration Guidelines for Aggregated Ethernet Interfaces

- Simple filters are not supported for interfaces in aggregated Ethernet bundles:
  - On M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
  - On MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

For more information about simple filters, see the *Junos OS Class of Service Library for Routing Devices*.

- On the aggregated Ethernet bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific capabilities, see *Gigabit Ethernet Accounting and Policing Overview*.
- Aggregated Ethernet interfaces can be either tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on MX Series routers support the configuration of **flexible-vlan-tagging** and **native-vlan-id** on dual-tagged frames, which consist of the following configuration statements:
  - inner-tag-protocol-id
  - inner-vlan-id
  - pop-pop
  - pop-swap
  - push-push
  - swap-push
  - swap-swap

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

Table 3 on page 4 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

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

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

Related Documentation

- inner-tag-protocol-id on page 227
- inner-vlan-id on page 228
- pop-pop on page 246
- pop-swap on page 247
- push-push on page 248
- swap-push on page 250
- swap-swap on page 251
- Configuring Mixed Aggregated Ethernet Links on page 41

- *Gigabit Ethernet Accounting and Policing Overview*
- *Ethernet Interfaces*

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

MX Series routers support active-active bridging and virtual router redundancy protocol (VRRP) over Integrated routing and bridging (IRB). This is a common scenario used in data centers. This section provides an overview of the supported functionality.

Active-active bridging and VRRP over IRB support extends multichassis link aggregation group (MC-LAG) by adding the following functionality:

- Interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for active-active bridging
- Active-active bridging
- VRRP over IRB for active-active bridging
- A single bridge domain cannot correspond to two redundancy group IDs

The following functionalities are supported for MC-LAG in an active-active bridging domain:

- MC-LAG is supported only between two chassis. Interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for active-active bridging Active-active bridging VRRP over IRB for active-active bridging.
- For VPLS networks, you can configure the aggregated Ethernet (aeX) interfaces on MC-LAG devices with the **encapsulation ethernet-vpls** statement to use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values or the **encapsulation vlan-vpls** statement to use Ethernet VLAN encapsulation on VPLS circuits.
- Layer 2 circuit functionalities are supported with **ethernet-ccc** as the encapsulation mode.
- Network topologies in a triangular and square pattern are supported. In a triangular network design, with equal-cost paths to all redundant nodes, slower, timer-based convergence can possibly be prevented. Instead of indirect neighbor or route loss detection using hellos and dead timers, you can identify the physical link loss and denote a path as unusable and reroute all traffic to the alternate equal-cost path. In a square network design, depending on the location of the failure, the routing protocol might converge to identify a new path to the subnet or the VLAN, causing the convergence of the network to be slower.
- Interoperation of Link Aggregation Control Protocol (LACP) for MC-LAG devices. LACP is one method of bundling several physical interfaces to form one logical interface. When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive

them from another link. One side of the link must be configured as active for the link to be up.

- Active-standby mode is supported using LACP. When an MC-LAG operates in the active-standby mode, one of the router's ports only becomes active when failure is detected in the active links. In this mode, the PE routers perform an election to determine the active and standby routers.
- Configuration of the pseudowire status type length variable (TLV) is supported. The pseudowire status TLV is used to communicate the status of a pseudowire back and forth between two provider edge (PE) routers. The pseudowire status negotiation process assures that a PE router reverts back to the label withdraw method for pseudowire status if its remote PE router neighbor does not support the pseudowire status TLV.
- The MC-LAG devices use Interchassis Communication Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

Keep the following points in mind when you configure MC-LAG in an active-active bridging domain:

- A single bridge domain cannot be associated with two redundancy groups. You cannot configure a bridge domain to contain logical interfaces from two different multichassis aggregated Ethernet (MC-AE) interfaces and associate them with different redundancy group IDs by using the redundancy group group-id statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level.
- You must configure logical interfaces in a bridge domain from a single MC-AE interface and associate it with a redundancy group. You must configure a service ID by including the **service-id *vid*** statement at the **[edit bridge-domains *bd-name*]** hierarchy level for MC-AE interfaces if you configure logical interfaces on MC-AE interfaces that are part of the bridge domain.

For a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain and with VRRP configured over an integrated routing and bridging (IRB) interface, you must include the **accept-data** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other layer 3 protocols and applications to work properly over multi-chassis aggregated Ethernet (mc-aeX) interfaces.

NOTE: On an MC-LAG, the unit number associated with aggregated Ethernet interfaces on provider Edge Router, PE1, must match the unit number associated with aggregated Ethernet interfaces on provider Edge Router, PE2. If the unit numbers differ, MAC address synchronization does not happen. As a result, the status of the MAC address on the remote provider edge router remains in pending state.

You can configure the **vlan-id none** statement at the **[edit bridge-domain *bd-name*]** hierarchy level for MC-LAG in an active-active bridge domain for Layer 2 networks and IPv4 networks. Starting with Junos OS Release 13.2R3, you can configure the **vlan-id none** statement at the **[edit bridge-domain *bd-name*]** hierarchy level for MC-LAG in an active-active bridge domain for IPv6 networks. The **vlan-id none** statement removes the incoming VLAN tags that identify a Layer 2 logical interface when packets are transmitted over VPLS pseudowires.

The topologies shown in Figure 1 on page 10 and Figure 2 on page 10 are supported. These figures use the following abbreviations:

- Aggregated Ethernet (AE)
- Interchassis link (ICL)
- Multichassis link (MCL)

Figure 1: Single Multichassis Link

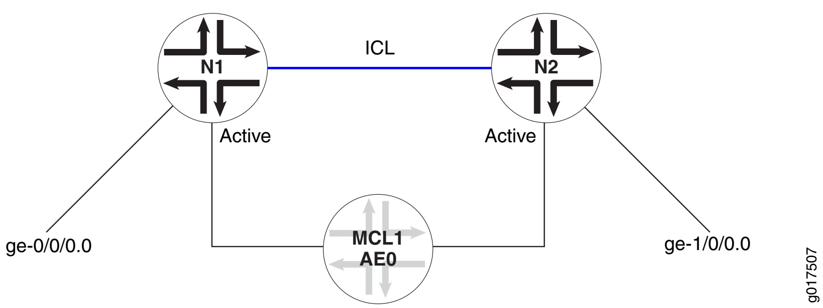

Figure 2: Dual Multichassis Link

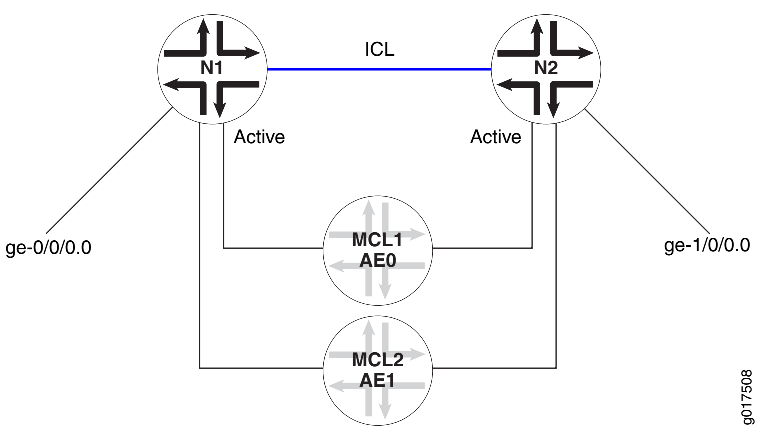

The following functionality is not supported:

- Virtual private LAN service (VPLS) within the core
- Bridged core
- Topology as described in Rule 4 of "Data Traffic Forwarding Rules" on page 13
- Routed multichassis aggregated Ethernet (RMC-AE), where the VRRP backup master is used in the edge of the network
- Track object, where in the case of an MC-LAG, the status of the uplinks from the provider edge can be monitored and the MC-LAG can act on the status
- Mixed mode (active-active MC-LAG is supported on MX series routers with MPC/MIC interfaces only). All interfaces in the bridge-domain that are mc-ae active-active, must be on MPC/MICs.

The topologies shown in Figure 3 on page 11, Figure 4 on page 11 and Figure 5 on page 12 are not supported:

Figure 3: Interchassis Data Link Between Active-Active Nodes

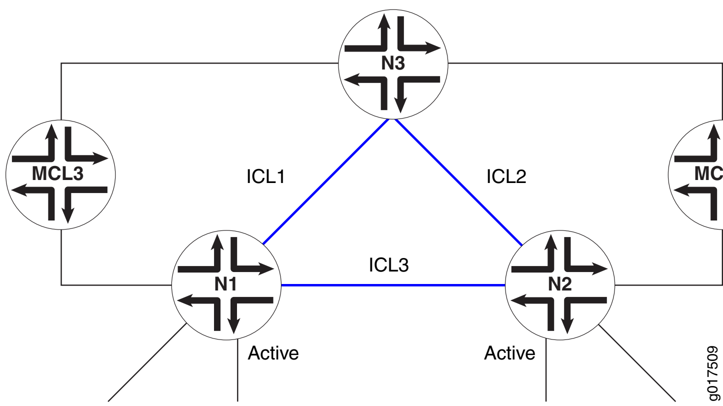

Figure 4: Active-Active MC-LAG with Single MC-LAG

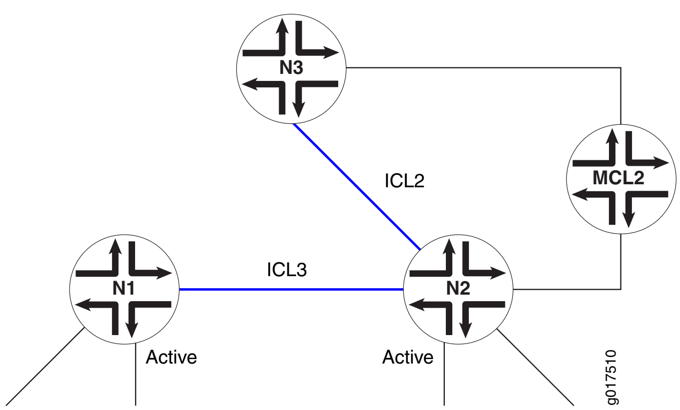

Figure 5: Active-Active MC-LAG with Multiple Nodes on a Single Multichassis Link

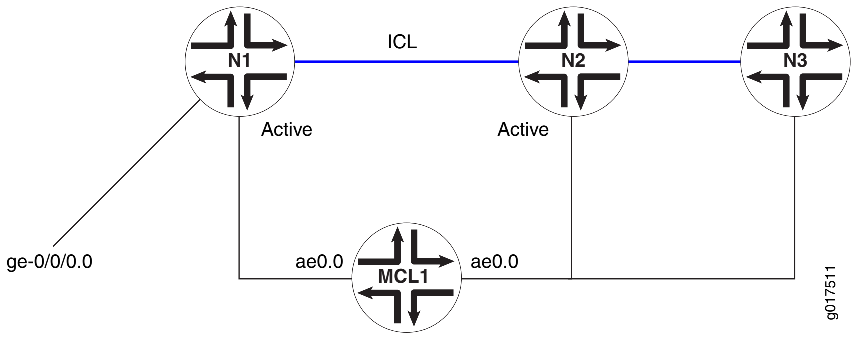

NOTE: A redundancy group cannot span more than two routers.

When configured to be active-active, the client device load balances the traffic to the peering MC-LAG network devices. In a bridging environment, this could potentially cause the following problems:

- Traffic received on the MC-LAG from one MC-LAG network device could be looped back to the same MC-LAG on the other MC-LAG network device.
- Duplicated packets could be received by the MC-LAG client device.
- Traffic could be unnecessarily forwarded on the interchassis link.

To better illustrate the problems listed above, consider Figure 6 on page 13, where an MC-LAG device MCL1 and single-homed clients **ge-0/0/0.0** and **ge-1/0/0.0** are allowed to talk to each other through an ICL:

- Traffic received on network routing instance N1 from MCL1 could be flooded to ICL to reach network routing instance N2. Once it reaches network routing instance N2, it could be flooded back to MCL1.
- Traffic received on interface **ge-0/0/0.0** could be flooded to MCL1 and ICL on network routing instance N1. Once network routing instance N2 receives such traffic from ICL, it could be again flooded to MCL1.
- If interface **ge-1/0/0.0** does not exist on network routing instance N2, traffic received from interface **ge-0/0/0.0** or MCL1 on network routing instance N1 could be flooded to network routing instance N2 through ICL unnecessarily since interface **ge-0/0/0.0** and MCL1 could reach each other through network routing instance N1.

Figure 6: MC-LAG Device and Single-Homed Client

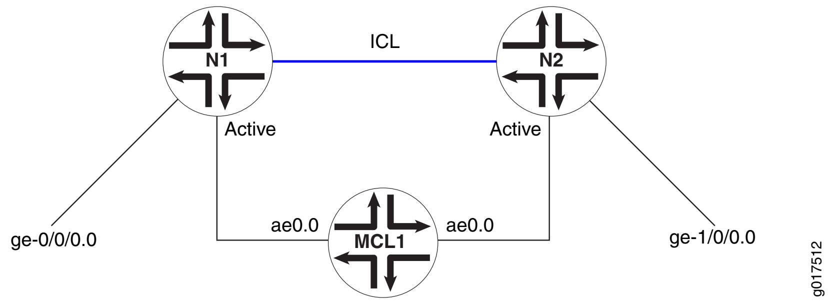

Advantages of Using Multichassis Link Aggregation Groups

An MC-LAG reduces operational expenses by providing active-active links with a LAG, eliminates the need for the Spanning Tree Protocol (STP), and provides faster layer 2 convergence upon link and device failures.

An MC-LAG adds node-level redundancy to the normal link-level redundancy that a LAG provides. An MC-LAG improves network resiliency, which reduces network down time as well as expenses.

In data centers, it is desirable for servers to have redundant connections to the network. You probably want active-active connections along with links from any server to at least two separate switches.

An MC-LAG allows you to bond two or more physical links into a logical link between two switches or between a server and a switch, which improves network efficiency. An MC-LAG enables you to load balance traffic on multiple physical links. If a link fails, the traffic can be forwarded through the other available link and the logical aggregated link remains in the UP state.

Data Traffic Forwarding Rules

In active-active bridging and VRRP over IRB topographies, network interfaces are categorized into three different interface types, as follows:

S-Links—Single-homed link (S-Link) terminating on MC-LAG-N device or MC-LAG in active-standby mode. In Figure 6 on page 13, interfaces **ge-0/0/0.0** and **ge-1/0/0.0** are S-Links.

MC-Links—MC-LAG links. In Figure 6 on page 13, interface **ae0.0** is the MC-Link.

ICL—Interchassis data link.

Based on incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations, as described in the data traffic forwarding rules. Note that if only one of the MC-LAG member link is in the UP state, it is considered an S-Link.

The following data traffic forwarding rules apply:

1. When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.

2. NOTE: **The functionality described in rule 2 is not supported.**

   The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:

   a. If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device.

   b. Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other.

   Only if both a. and b. are true, is traffic always forwarded to the ICL.

3. When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network that the packet comes from.

4. NOTE: **The topology shown in Figure 7 on page 15 is not supported.**

   In certain cases, for example the topology shown in Figure 7 on page 15, there could be a loop caused by the ICL. To break the loops, one of the following mechanisms could be used:

   a. Run certain protocols, such as spanning tree protocol (STP). In this case, whether packets received on one ICL are forwarded to other ICLs is determined by using Rule 3.

   b. Configure the ICL to be fully meshed among the MC-LAG network devices. In this case, traffic received on the ICL would be not be forwarded to any other ICLs.

   In either case, duplicate packets could be forwarded to the MC-LAG clients. Consider the topology shown in Figure 7 on page 15, where if network routing instance N1 receives a packet from **ge-0/0/0.0**, it could be flooded to ICL1 and ICL3.

   When receiving from ICL1 and ICL3, network routing instances N3 and N2 could flood the same packet to MCL2, as shown in Figure 7 on page 15. To prevent this from happening, the ICL designated forwarder should be elected between MC-LAG peers and traffic received on an ICL could be forwarded to the active-active MC-LAG client by the designated forwarder only.

Figure 7: Loop Caused by the ICL Links

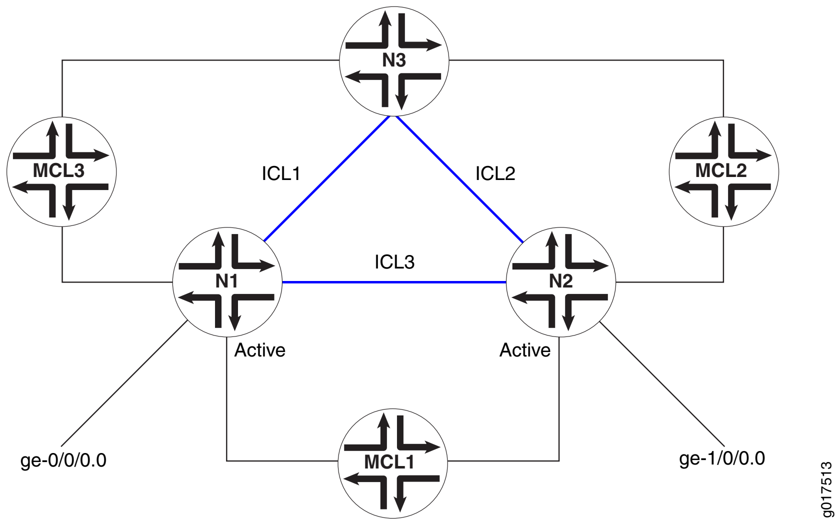

5. When received from an ICL, traffic should not be forwarded to the core-facing client link connection between two provider edge (PE) devices (C-Link) if the peer chassis's (where the traffic is coming from) C-Link is UP.

MAC Address Management

If an MC-LAG is configured to be active-active, upstream and downstream traffic could go through different MC-LAG network devices. Since the media access control address (MAC address) is learned only on one of the MC-LAG network devices, the reverse direction's traffic could be going through the other MC-LAG network and flooded unnecessarily. Also, a single-homed client's MAC address is only learned on the MC-LAG network device it is attached to. If a client attached to the peer MC-LAG network needs to communicate with that single-homed client, then traffic would be flooded on the peer MC-LAG network device. To avoid unnecessary flooding, whenever a MAC address is learned on one of the MC-LAG network devices, it gets replicated to the peer MC-LAG network device. The following conditions should be applied when MAC address replication is performed:

- MAC addresses learned on a MC-LAG of one MC-LAG network device should be replicated as learned on the same MC-LAG of the peer MC-LAG network device.
- MAC addresses learned on single-homed customer edge (CE) clients of one MC-LAG network device should be replicated as learned on ICL-PL interface of the peer MC-LAG network device.
- MAC addresses learned on MC-LAG VE clients of one MC-LAG network device should be replicated as learned on the corresponding VE interface of the peer MC-LAG network device.
- MAC address learning on an ICL is disabled from the data path. It depends on software to install MAC addresses replicated through interchassis control protocol (ICCP).

MAC Aging

MAC aging support in the Junos OS extends aggregated Ethernet logic for a specified MC-LAG. A MAC address in software is deleted until all Packet Forwarding Engines have deleted the MAC address. In the case of an MC-LAG, a remote provider edge is treated as a remote Packet Forwarding Engine and has a bit in the MAC data structure.

Layer 3 Routing

In general, when an MC-LAG is configured to provide Layer 3 routing functions to downstream clients, the MC-LAG network peers should be configured to provide the same gateway address to the downstream clients. To the upstream routers, the MC-LAG network peers could be viewed as either equal-cost multi path (ECMP) or two routes with different preference values.

Junos OS supports active-active MC-LAGs by using VRRP over IRB. Junos OS also supports active-active MC-LAGs by using IRB MAC address synchronization. You must configure IRB using the same IP address across MC-LAG peers. IRB MAC synchronization is supported on 32-bit interfaces and interoperates with earlier MPC/MIC releases.

To ensure that Layer 3 operates properly, instead of dropping the Layer 3 packet, the VRRP slave attempts to perform routing functions if the packet is received on an MC-LAG. A VRRP slave sends and responds to address resolution protocol (ARP) requests.

For ARP, the same issue exists as with Layer 2 MAC addresses. Once ARP is learned, it must be replicated to the MC-LAG through ICCP. The peer must install an ARP route based on the ARP information received through ICCP.

For ARP aging, ARP requests on the MC-LAG peers can be aged out independently.

Address Resolution Protocol Active-Active MC-LAG Support Methodology

Suppose one of the PE routers issues an ARP request and another PE router gets the response and, because of the aggregated Ethernet distribution logic, the ARP resolution is not successful. Junos OS uses ARP response packet snooping to perform active-active multichassis link aggregation group support, providing easy synchronization without the need to maintain any specific state.

IGMP Snooping on Active-Active MC-LAG

For multicast to work in an active-active MC-LAG scenario, the typical topology is as shown in Figure 8 on page 17 and Figure 9 on page 18 with interested receivers over S-links and MC-Links. Starting in Junos OS Release 11.2, support is extended for sources connected over Layer 2 interface.

If an MC-LAG is configured to be active-active, reports from MC-LAG clients could reach any of the MC-LAG network device peers. Therefore the IGMP snooping module needs to replicate the states such that the Layer 2 multicast route state on both peers are the same. Additionally for S-Link clients, snooping needs to replicate these joins to its snooping peer, which in the case of Layer 3 connected source, passes this information to the PIM on IRB to enable the designated router to pull traffic for these groups, The ICL should be configured as a router facing interface. For the scenario where traffic arrives via a Layer 3 interface, it is a requirement to have PIM and IGMP enabled on the IRB interface configured on the MC-LAG network device peers.

Figure 8: Multicast Topology with Source Connected via Layer 3

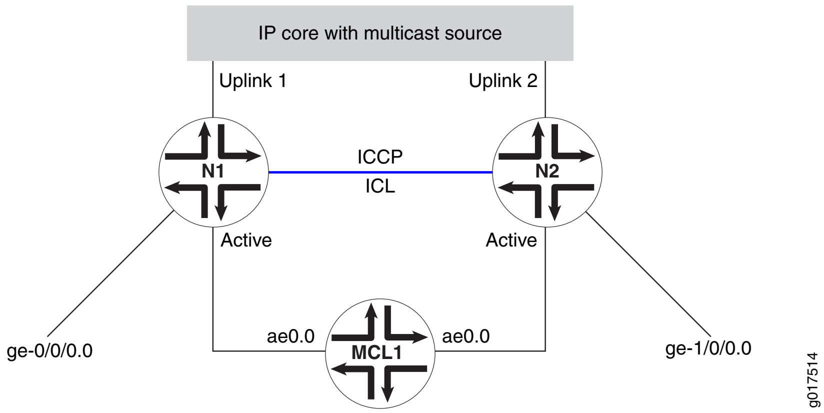

With reference to Figure 8 on page 17, either N1 or N2 becomes a designated router (for this example, N1 is the designated router). Router N1 would therefore pull the multicast traffic from the core. Once multicast data hits the network device N1, the data is forwarded based on the snooping learned route.

For MC-Link clients, data is forwarded via N1. In the case of failover of the MC-Links, the data reaches the client via N2. For S-Link clients on N1, data would be forwarded via normal snooping routes.

For S-Link clients on N2, data is forwarded via the ICL interface. Layer 2 multicast routes on N1 do not show these groups unless there is interest for the same group over MC-Links or over S-Links on N1. For IRB scenario, the IGMP membership and Layer 3 multicast route on N1 does however show these groups learned over the IRB interface.

Therefore, for a case where a specific group interest is only on the S-Link on N2, data arriving on N1 reaches N2 via the default route and the Layer 2 multicast route on N2 has the S-Link in the outgoing interface list.

Figure 9: Multicast Topology with Source Connected via MC-Link

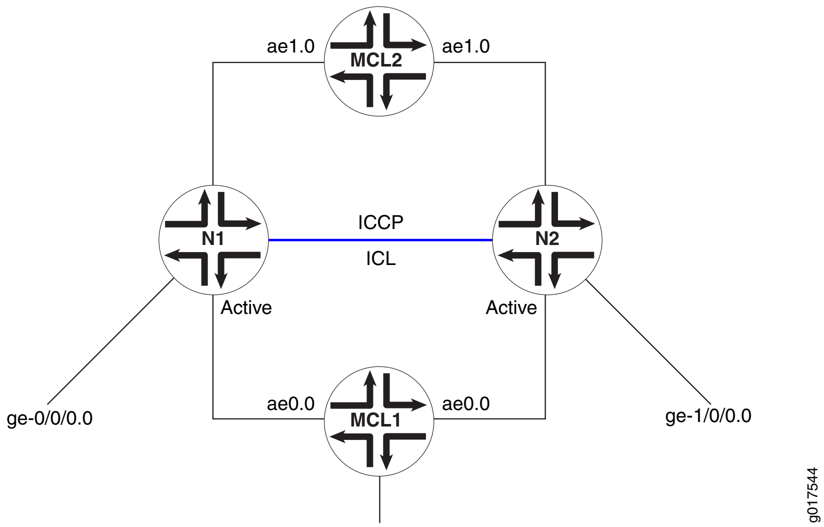

In Figure 9 on page 18, MCL1 and MCL2 are on different devices and the multicast source or IGMP querier is connected via MCL2. The data forwarding behavior seen is similar to that explained for multicast topology with source connected via Layer 3.

NOTE: **IGMP snooping should not be configured in proxy mode. There should be no IGMP hosts or IGMP/PIM routers sitting on the ICL interface.**

Up and Down Event Handling

The following conditions apply to up and down event handling:

1. If the interchassis control protocol (ICCP) connection is UP but the ICL interface becomes DOWN, the router configured as standby will bring down all the multichassis aggregated Ethernet interfaces shared with the peer which is connected to ICL. This will make sure that there are no loops in the network. Otherwise, both PEs will becomes PIM designated routers and, hence, forward multiple copies of the same packet to the customer edge.
2. If the ICCP connection is UP and the ICL comes UP, the router configured as standby will bring up the multichassis aggregated Ethernet interfaces shared with the peer.
3. If both the ICCP connection and the ICL are DOWN, the router configured as standby will bring up the multichassis aggregated Ethernet interfaces shared with the peer.
4. The layer 2 address learn daemon (l2ald) does not store the information about a MAC address learned from a peer in the kernel. If l2ald restarts, and if the MAC address was not learned from the local multichassis aggregated Ethernet interface, l2ald will clear the MAC addresses and this will cause the router to flood the packets destined to this MAC address. This behavior is similar to that in a Routing Engine switchover. (Please note that currently l2ald runs on a Routing Engine only when it is a master). Also, during the time l2ald is DOWN, ARP packets received from an ICCP peer will be

dropped. ARP retry will take care of this situation. This will be the case with Routing Engine switchover too.

5. If ICCP restarts, l2ald does not identify that a MAC address was learned from a peer and, if the MAC address was learned only from the peer, that MAC address will be deleted and the packets destined to this MAC address will be flooded.

VRRP Active-Standby Support

VRRP in active-standby mode enables Layer 3 routing over the MC-AE interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The virtual routers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address, known as a routed VLAN interface (RVI), maps to either of the VRRP MAC addresses or the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets. At any time, one of the VRRP routers is the master (active), and the other is a backup (standby). Both VRRP active and VRRP backup routers forward Layer 3 traffic arriving on the MC-AE interface. If the master router fails, all the traffic shifts to the MC-AE link on the backup router.

NOTE: **You must configure VRRP on both MC-LAG peers in order for both the active and standby members to accept and route packets. Additionally, configure the VRRP backup router to send and receive ARP requests.**

Routing protocols run on the primary IP address of the RVI, and both of the MC-LAG peers run routing protocols independently. The routing protocols use the primary IP address of the RVI and the RVI MAC address to communicate with the MC-LAG peers. The RVI MAC address of each MC-LAG peer is replicated on the other MC-LAG peer and is installed as a MAC address that has been learned on the ICL-PL.

In many cases, MC-LAG devices are Layer 3 routing devices and perform the default gateway functionality for hosts that are part of the attached IP subnets. The MC-LAG device-pair can therefore share the default gateway routing functionality with the VRRP protocol. The VRRP active-standby operation can be optimized to be an active-active mode of processing because traffic flowing from an MC-LAG client to an MC-LAG device is always sent on one of the available links of an MC-LAG and reaches exactly one of the MC-LAG destination devices. Because of this behavior, both MC-LAG devices in the pair can be enabled as routers for IP traffic destined to the VRRP destination MAC address. Both MC-LAG devices must be full members of the routing domain and have routing entries that allow them to reach the IP destination networks.

The active-active functionality works without changing any VRRP state machine and by activating the routing function in the forwarding plane of the VRRP backup system (similar to the VRRP master). The mechanism enables traffic forwarding in Layer 3 to be fully redundant, leveraging all available link bandwidth. All routers forward traffic and thereby load share routed traffic.

Interchassis Control Protocol

Interchassis control protocol (ICCP) is used to synchronize configurations, states, and data.

ICCP supports the following types of state information:

- MC-LAG members and their operational states.
- Single-homed members and their operational states.

ICCP supports the following application database synchronization parameters:

- MAC addresses learned and to be aged.
- ARP info learned over IRB.

Interchassis Control Protocol Message

ICCP messages and attribute-value pairs (AVPs) are used for synchronizing MAC address and ARP information.

Related Documentation

- Configuring Multichassis Link Aggregation on page 42
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on page 127
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 56
- **multi-chassis-protection on page 242**
- **peer on page 244**
- **show interfaces mc-ae on page 378**
- *Ethernet Interfaces*

IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview

IGMP Snooping in MC-LAG Active-Active on MX Series Routers Functionality

MX Series routers support multichassis link aggregation group (MC-LAG) active-active and IGMP snooping in active-standby mode. MC-LAG allows one device to form a logical LAG interface with two or more network devices. MC-LAG provides additional benefits including node level redundancy, multi-homing, and loop-free layer-2 network without running STP. The following features are supported:

- State synchronization between peers for IGMP snooping in a bridge domain with only Layer 2 interfaces
- Qualified learning
- Router facing multichassis links

MX Series routers support the following enhancements to active-active bridging and virtual router redundancy protocol (VRRP) over integrated routing and bridging (IRB):

- MC-LAG support for IGMP snooping in a pure Layer 2 switch
- MC-LAG support for IGMP snooping in bridge domains doing qualified learning
- Support for MC-Links being router facing interfaces

The following functions are not supported:

- MC-LAG for VPLS instances
- MC-Links trunk ports
- Proxy mode for active-active
- Adding interchassis links to outgoing interfaces on an as needed basis. Interchassis links can be added to the outgoing interface list as router facing interfaces.

Typically Supported Network Topology for IGMP Snooping with MC-LAG Active-Active Bridging

Figure 10 on page 22 depicts a typical network topology over which IGMP snooping with MC-LAG active-active is supported.

Figure 10: Typical Network Over Which Active-Active Is Supported

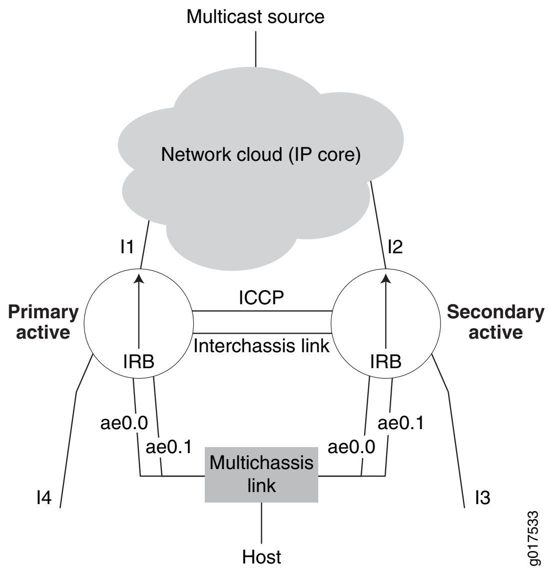

Interfaces I3 and I4 are single-homed interfaces. The multichassis links (MC-Link) ae0.0 and ae0.1 belong to the same bridge domain in both the chassis. Interfaces I3, ae0.0 and ae0.1 are in the same bridge domain in S-A. Interfaces I4, ae0.0 and ae0.1 are in the same bridge domain in the primary active (P-A) router. Interfaces I3, I4, ae0.0 and ae0.1 are in the same learning domain as is the interchassis link (ICL) connecting the two chassis.

The primary active router is the chassis in which the integrated routing and bridging has become PIM-DR. The secondary active router is the chassis in which integrated routing and bridging is not PIM DR. Router P-A is the chassis responsible for pulling traffic from the IP core. Hence, PIM-DR election is used to avoid duplication of data traffic.

Learning domains are described in “Qualified Learning” on page 24.

For the IGMP speakers (hosts and routers) in the learning domain, P-A and S-A together should appear as one device with interfaces I4, I3, ae0.0 and ae0.1.

No duplicate control packets should be sent on multichassis links, meaning the control packet should be sent through only one link.

Control Plane State Updates Triggered by Packets Received on Remote Chassis

The membership state in Layer 3 multicast routing is updated as a result of reports learned on remote legs of multichassis links and s-links attached to the remote chassis.

The membership state and routing entry in snooping is updated when reports are received on the remote legs of a multichassis link.

When reports are received on S-links attached to the remote chassis the membership state or routing entry in snooping is not updated.

The list of <s,g>s for which the state is maintained is the same in both the chassis under snooping as long as the outgoing interface lists involve only multichassis links.

Data Forwarding

This discussion assumes integrated routing and bridging on P-A is the PIM-DR. It pulls the traffic from sources in the core. Traffic might also come on Layer 2 interfaces in the bridge domain. For hosts directly connected to the P-A chassis, there is no change in the way data is delivered.

For delivering traffic to hosts connected to S-A (which is the non-DR) on the single-homed link like l3, we rely on interchassis link. The traffic that hits P-A is sent over ICL to S-A to be delivered to the links that have reported interests in s,g and the links that are router facing.

When ae0 leg in P-A goes down, the hosts connected to the multichassis link will receive traffic via ICL. In S-A, traffic received on ICL is sent to multichassis links in the outgoing interface list for which the ae counterpart in P-A is down.

Pure Layer 2 Topology Without Integrated Routing and Bridging

Figure 11 on page 24 illustrates the chassis connecting to the PIM-DR is the primary active router and the other is the secondary active.

Figure 11: Layer 2 Configuration Without Integrated Routing and Bridging

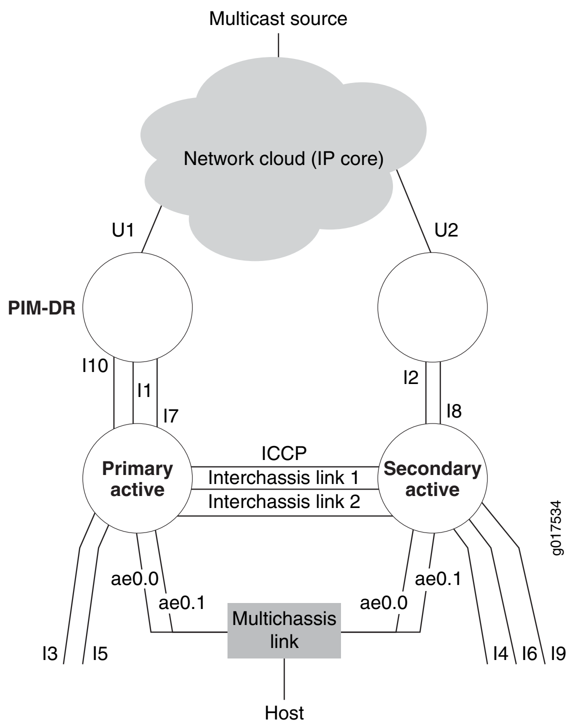

Qualified Learning

In this application, interfaces I1, I2, I3, I4, I5, I6, I7, I8, I9 and I10 are single-homed interfaces. The multichassis links ae0.0 and ae0.1 belong to the same bridge domain in both the chassis. Interfaces I10, I1,I7,I3,I5, ae0.0 and ae0.1 are in same bridge domain, bd1 in P-A. Interfaces I9, I2,I8,I4,I6, ae0.0 and ae0.1 are in same bridge domain, bd1 in S-A.

This discussion assumes the following configuration:

- In Primary Active and S-A, qualified learning is ON in bd1.
- Interfaces I1,I2,I3, ae0.0 and I4 belong to vlan1, learning domain ld1.
- Interfaces I7,I8,I5, ae0.1 and I6 belong to vlan2, learning domain ld2.
- Interfaces I9 and I10 belong to vlan3, learning domain ld3.

For the IGMP speakers (hosts and routers) in the same learning domain ld1, P-A and S-A linked should appear to be one switch.

For the IGMP speakers (hosts and routers) in the same learning domain ld2, P-A and S-A linked should appear to be one switch.

Since there are no multichassis links in learning domain ld3, for the IGMP speakers (hosts and routers) in learning domain ld3, P-A and S-A will not appear to be one switch.

This discussion assumes interchassis link ICL1 corresponds to learning domain ld1 and interchassis link ICL2 corresponds to learning domain ld2.

Control packet flow is supported, with the exception of passing information to IRB.

Data Forwarding with Qualified Learning

This discussion assumes one learning domain (LD), ld1, and further assumes interface I1 on router P-A is connected to the PIM-DR in the learning domain and pulls the traffic from sources in the core.

For delivering traffic to hosts connected to router S-A (which is the non-DR) on the single-homed link like I2, I4 (belonging to ld1), we rely on ICL1. The traffic that hits router P-A on interface I1 is sent over interchassis link ICL1 to router S-A to be delivered to the links that have reported interests in s,g or the links that are router facing in learning domain ld1.

When the interface ae0 leg in router P-A goes down, the hosts connected to the multichassis link receive traffic from interface I1 via the interchassis link ICL1. In router S-A, traffic received on interchassis link ICL1 is sent to multichassis links in the outgoing interface list for which the aggregated Ethernet counterpart in router P-A is down.

It is further assumed that interface I9 in router S-A belongs to the learning domain ld3 with interests in s,g, and that interface I10 in learning domain ld3 in router P-A receives traffic for s,g. Interface I9 does not receive data in this topology because there are no multichassis links (in a-a mode) and hence no interchassis link in learning domain ld3.

Static Groups on Single Homed Interfaces

For multichassis links, the static group configuration should exist on both legs and synchronization with the other chassis is not required.

Synchronization of the static groups on single homed interfaces between the chassis is not supported, however the addition of logical interfaces to the default outgoing interface list supports traffic delivery to the interface within a static configuration.

Router Facing Interfaces as Multichassis Links

IGMP queries could arrive on either leg of the multichassis links but in both peers, the multichassis link should be considered as router facing.

Reports should exit only once from the multichassis link, that is from only one leg.

The following MC-LAG support for IGMP snooping in IRB is provided:

- Non-proxy snooping
- Logical interfaces must be outgoing interfaces for all routes including the default route
- IGMP snooping in a pure Layer 2 switch

- IGMP snooping in bridge domains doing qualified learning
- Router facing interface MC-Links

The following features are not supported:

- Proxy mode for active-active
- MC-LAG support for VPLS instances
- Trunk ports as multichassis links
- Adding logical interfaces to outgoing interfaces on need basis. However, logical interfaces are always added as a router facing interface to the outgoing interface list.

Related Documentation

- Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers on page 130
- Example: Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers on page 131
- *Example: Configuring IGMP Snooping*
- **igmp-snooping on page 226**
- **multicast-router-interface on page 241**
- **show l2-learning instance on page 364**
- *Ethernet Interfaces*

Understanding Aggregated Ethernet Load Balancing

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

  Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

  By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

  ECMP load balancing can be:

  - Across BGP paths (BGP multipath)
  - Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

  Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

  To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  To configure the tolerance value as a percentage, include the **tolerance** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the **pps** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the **scan-interval** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

NOTE: The **pps** and **scan-interval** optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-Packet Random Spray Load Balancing

  When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

NOTE: The Per-Packet option for load balancing is not supported on PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the **show interfaces ae*x* aggregated-ether-options load-balance** command.

Related Documentation

- Example: Configuring Aggregated Ethernet Load Balancing on page 151

Multichassis Link Aggregation on Logical Systems Overview

On MX Series routers, multichassis link aggregation (MC-LAG) enables a device to form a logical LAG interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node level redundancy, multi- homing support, and loop-free Layer 2 network without running Spanning Tree Protocol (STP). The MC-LAG devices use Inter-Chassis Communication Protocol (ICCP) to exchange the control information between two MC-LAG network devices. Starting with Junos OS Release 14.1, you can configure MC-LAG interfaces on logical systems within a router.

To configure ICCP for MC-LAG interfaces on logical systems, include the **iccp** statement at the **[edit logical-systems *logical-system-name* protocols]** hierarchy level. To view ICCP information for MC-LAG on logical systems, use the **show iccp logical-system *logical-system-name*** command. To view ARP statistics or remote MAC addresses for the multichassis aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a logical system, use the **show l2-learning redundancy-groups *group-name* logical-system *logical-system-name* (arp-statistics | remote-macs)** command. To view Neighbor Discovery (ND) statistical details for MC-AE nodes on redundancy groups of a logical group, use the **show l2-learning redundancy-groups *group-name* logical-system *logical-system-name* nd-statistics** command.

Logical systems enable effective, optimal segregation of a single router into multiple virtual partitions, which can be configured and managed by diversified entities. Logical systems perform a subset of the actions of a physical router and have their own unique routing tables, interfaces, policies, and routing instances. A set of logical systems within a single router can handle the functions previously performed by several small routers.

In a network deployment that contains MC-LAG interfaces, you can configure such interfaces on logical systems contained within a router. When you configure multichassis aggregated Ethernet (MC-AE) interfaces on a logical system, you must ensure that these interfaces are added with the same multichassis aggregated Ethernet identification number and redundancy group identifier for the MC-LAG on both the peers or devices that are connected by the MC-AE interfaces. It is not necessary to specify the same logical system name on both the peers; however, you must ensure that the Inter-Chassis Control Protocol (ICCP) to associate the routing or switching devices contained in a redundancy group is defined on both the peers within the logical systems of the devices. Such a configuration ensures that all the packets are transmitted using ICCP within the logical sysem network. The logical system information is added and removed by the ICCP process to prevent each packet from containing the logical system details. This behavior enables multiple disjoint users to employ MC-LAG capabilities within their networks transparently and seamlessly. A unique ICCP definition for a logical system is created, thereby enabling you to wholly manage the ICCP parameters on one logical system without the need for access permissions to view other logical system networks on the same device. Configuration of MC-LAG interfaces on logical systems enables MC-LAG to be used across multiple routing tables and switch forwarding tables in active-active and active-standby modes of MC-LAG interfaces.

Because the Layer 2 address learning process (L2ALD) supports logical systems, the ARP, neighbor discovery (ND), and MAC synchronization packets that are traversing an MC-AE interface use the logical system:routing instance (LS:RI) combination to map the packets to the correct routing instance in a logical system. LACP does not require the LS-RI combination to be identified because it operates on physical interfaces and is unique within a chassis. For a service in the set of provider edge (PE) routers providing the service, the service ID distinguishes the routing instances in a logical system because it is unique for a logical system across a routing instance. MC-LAG is configured on the aggregated Ethernet (ae-) bundle interface. An ae- interface is a logical interface and is globally unique, which causes the MCLAG configuration to be exclusive and separate for a router. You can add ae- interfaces in an MC-LAG configuration to be part of a logical system and use it throughout that particular logical system.

Sample Configuration Scenario for MC-LAG on Logical Systems

Consider a sample scenario in which two MX Series routers, MX1 and MX2, are connected using an ae- interface that is enabled with MC-LAG. The peers in an MC-LAG use an interchassis control link-protection link (ICL-PL) to replicate forwarding information across the peers. Additionally, ICCP propagates the operational state of MC-LAG members through the ICL-PL. The two PE devices, MX1 and MX2, each have a LAG connected to the CE devices, CE1 and CE2. Four logical systems are defined on each of the PE devices, MX1 and MX2. CE-1 and CE-2 can be part of the same VLAN with the same VLAN ID and located in the same IP subnet for MC-LAG in two different Logical systems. All the four logical system entities can work independently in MX1 and MX2.

The ICCP process can manage multiple client-server connections with its peer ICCP instances based on the ICCP configuration for the logical system:routing instance (LS-RI) combinations. Each ICCP connection is associated with an LS-RI combination. For example, with two routing instances, IP1 and IP2 on each of the logical systems, LS1 and LS2, the following mapping is performed for ICCP settings:

[ICCP] (LS1) (IP1) < = = > (IP2) (LS1) [ICCP] within LS1 network.

[ICCP] (LS2) (IP1) < = = > (IP2) (LS2) [ICCP] within LS2 network.

An ICCP instance in a logical system is linked with the ICCP instance of the peer logical system. The ICCP application transmits the relevant routing index depending on the LS:RI combination to the BFD process, when BFD is configured in your topology.

The L2ALD process transmits and receives ARP, neighbor discovery, and MAC synchronization packets with the LS-RI information. When the peer MAC synchronization packets are received, the L2ALD process decodes the LS details from the packet and determines whether an identical LS has been previously created on the router. If a match is found for the LS, the MAC forwarding entry for the corresponding bridge table for an interface bridge domain is created. If the LS in the received packet does not match with the defined LS on the device, for the MAC synchronization packet, the default logical instance is used for processing. Similarly, upon receipt of the ARP and ND packets, the L2ALD process decapsulates the LS information from packet and verifies if the corresponding logical instance has been previously created. If a match is found for the LS, the ARP and ND packets are processed according to the Layer 3 index that is unique in the system. Programming kernel entry mayn't require any LS info since it is programmed on L3 index which is unique in the system. If the LS in the received packet does not match with the defined LS on the device, for the ARP and ND packets, the default logical instance is used for processing. The routing instance is determined using the service ID attribute. The LS information is forwarded to ICCP, which in turn identifies the appropriate ICCP interface for the logical system and sends packets over it.

Guidelines for Configuring MC-LAG on Logical Systems

Keep the following points in mind while configuring MC-LAG interfaces on logical systems:

- You cannot use a single chassis to function as a provider edge (PE) and a customer edge (CE) device in different logical systems.
- You cannot use a single chassis to function as two PE devices by configuring logical systems on the chassis and ICCP, ICL links between the two logical systems because the MC-AE ID is unique in a router.
- IGMP snooping in MC-LAG topologies with logical systems is not supported.
- VPLS and VPN protocols with MC-LAG in active-standy mode is not supported.
- Logical system information is not communicated to the peer chassis because this detail is derived from an ICCP instance.

Related Documentation

- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on Logical Systems on MX Series Routers on page 72

PART 2

Configuration

CHAPTER 2

Aggregated Ethernet Interfaces

- Configuring Untagged Aggregated Ethernet Interfaces on page 172
- Configuring Aggregated Ethernet Link Speed on page 174
- Configuring Aggregated Ethernet Minimum Links on page 175
- Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces on page 176
- Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection on page 177
- Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers on page 178

Configuring an Aggregated Ethernet Interface

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

To configure an aggregated Ethernet interface:

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

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

2. Configure the aggregated Ethernet interface.

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

You specify the interface instance number ***x*** to complete the link association; ***x*** can be from 0 through 127, for a total of 128 aggregated interfaces on M Series and T Series routers and can be from 1 through 480, for a total of 480 aggregated interfaces on MX Series routers. You must also include a statement defining **ae*x*** at the **[edit interfaces]** hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see *Ethernet Interfaces Overview*, and for a sample configuration, see "Example: Configuring Aggregated Ethernet Interfaces" on page 35.

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

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

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

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

Related Documentation

- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 164
- Deleting an Aggregated Ethernet Interface on page 40
- Aggregated Ethernet Interfaces Overview on page 3

- *Ethernet Interfaces*

Example: Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count 15;
  }
}
```

```
[edit interfaces]
ge-1/3/0 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    link-speed 1g;
    minimum-links 1;
  }
  }
  vlan-tagging;
  unit 0 {
    vlan-id 1;
    family inet {
      address 14.0.100.50/24;
    }
  }
  unit 1 {
    vlan-id 1024;
    family inet {
      address 14.0.101.50/24;
```

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

Related Documentation

- *Ethernet Interfaces*
- Configure 'link-speed' for Gigabit Ethernet based Aggregate Ethernet interface bundles

Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a link aggregation group (LAG) or bundle. Link aggregation can be used for point-to-point connections. It balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

This topic describes how to configure aggregated Ethernet interfaces on PTX Series Packet Transport Routers.

On PTX Series Packet Transport Routers, aggregated Ethernet support includes the following features:

- A consistent interface type (**et** ***fpc/pic/port***) across all Ethernet interfaces.
- Ability to bundle multiple Ethernet interfaces
- Fault tolerance
- Load balancing between child links
- Advanced features including flexible VLAN tagging and Ethernet services encapsulation

Aggregated Ethernet interfaces can use interfaces from different FPCs or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
  aggregated-devices {
  ethernet {
    device-count 2;
  }
```



```
}
[edit interfaces]
et-0/0/0 {
  gigether-options {
    802.3ad ae0;
    }
  }
  et-0/0/1 {
    gigether-options {
      802.3ad ae0;
    }
  }
  ae0 {
    vlan-tagging;
    unit 0 {
      vlan-id 100;
      family inet {
        address 200.200.1.2/24;
      }
    }
    unit 1 {
      vlan-id 101;
      family inet {
        address 200.200.2.2/24;
      }
    }
  }
```

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 3
- Configuring Junos OS for Supporting Aggregated Devices on page 37

Configuring Junos OS for Supporting Aggregated Devices

Junos OS supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard.

Tasks for configuring aggregated devices are:

Configuring Virtual Links for Aggregated Devices

To define virtual links, you need to specify the associations between physical and logical devices within the **[edit interfaces]** hierarchy, and assign the correct number of logical devices by including the **device-count** statement at the **[edit chassis aggregated-devices ethernet]** and **[edit chassis aggregated-devices sonet]** hierarchy levels:

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count number;
  }
  sonet {
    device-count number;
  }
}
```

The maximum number of Ethernet logical interfaces that you can configure is 128. On M Series and T Series routers, you can configure a maximum number of 128 aggregated interfaces. On MX Series routers, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces are numbered from **ae0** through **ae127** for M Series and T Series routers, and the aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers. The maximum number of SONET/SDH logical interfaces is 64. The aggregated SONET/SDH interfaces are numbered from **as0** through **as63**.

NOTE: Starting with Junos OS Release 13.2, a maximum of 64 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces. In releases before Junos OS Release 13.2, a maximum of 16 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces.

Configuring LACP Link Protection at the Chassis Level

Link Aggregation Control Protocol (LACP) is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled. LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP link protection enables you to force active and standby links within an aggregated Ethernet. You configure LACP link protection by using the **link-protection** and **system-priority** statements at either the chassis or interface level and by configuring port priority at the interface level using the **system-priority** statement. Configuring LACP parameters at the chassis level results in all aggregated Ethernet interfaces using the defined values unless overridden by the LACP configuration on a specific interface.

```
[edit chassis]
aggregated-devices {
  ethernet {
    lacp {
      link-protection {
        non-revertive;
      }
      system-priority priority;
    }
  }
}
```


NOTE: LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

Enabling LACP Link Protection

To enable LACP link protection for aggregated Ethernet interfaces on the chassis, use the **link-protection** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
link-protection {
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, after a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch, and the current link remains active.

CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

Configuring System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the chassis, use the **system-priority** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
system-priority priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 through 65,535.

Configuring the Maximum Links Limit

To configure the maximum links limit, use the **maximum-links** statement at the **[edit chassis aggregated-devices]** hierarchy level:

```
[edit chassis aggregated-devices]
```

maximum-links *maximum-links-limit*;

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 34
- *Ethernet Interfaces*
- Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers on page 36
- *Configuring Aggregated SONET/SDH Interfaces*

Deleting an Aggregated Ethernet Interface

There are two approaches to deleting an aggregated Ethernet interface:

- You can delete an aggregated Ethernet interface from the interface configuration. The Junos OS removes the configuration statements related to **ae***x* and sets this interface to down state.
- You can also permanently remove the aggregated Ethernet interface from the device configuration by deleting it from the device-count on the routing device.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

   This step changes the interface state to down and removing the configuration statements related to **ae***x*.

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

2. Delete the interface from the device count.

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

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 34
- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 164
- Aggregated Ethernet Interfaces Overview on page 3
- *Ethernet Interfaces*

Configuring Mixed Aggregated Ethernet Links

In releases before Junos OS Release 13.2, all interfaces that form an aggregated Ethernet bundle must have the same speed and must be in full-duplex mode. Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

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

NOTE:

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

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

To configure mixed aggregated Ethernet interfaces:

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

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

   For example:

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

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

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

   For example:

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


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

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

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

For example:

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

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

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

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

For example:

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

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

```
[edit interfaces]
user@host# set aex unit number family inet address address
```

For example:

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

6. Commit the configuration:

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

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 3
- Configuring Aggregated Ethernet Link Speed on page 174
- link-speed on page 234

Configuring Multichassis Link Aggregation

On MX Series routers, multichassis link aggregation (MC-LAG) enables a device to form a logical LAG interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node level redundancy, multi-homing support, and loop-free Layer 2 network without running Spanning Tree Protocol (STP). MC-LAG can be configured for VPLS routing instance, CCC application, and Layer 2 circuit encapsulation types.

The MC-LAG devices use Inter-Chassis Communication Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

On one end of MC-LAG is a MC-LAG client device that has one or more physical links in a link aggregation group (LAG). This client device does not need to be aware of MC-LAG. On the other side of MC-LAG are two MC-LAG network devices. Each of these network devices has one or more physical links connected to a single client device. The network devices coordinate with each other to ensure that data traffic is forwarded properly.

MC-LAG includes the following functionality:

- Active standby mode is supported using Link Aggregation Control Protocol (LACP)
- MC-LAG operates only between two chassis.
- Layer 2 circuit functions are supported with **ether-ccc** and **vlan-ccc** encapsulation.
- VPLS functions are supported with **ether-vpls** and **vlan-vpls**.

NOTE: Ethernet connectivity fault management (CFM) specified in IEEE 802.1ag standard for Operation, Administration, and Management (OAM) is *not* supported on MC-LAG interfaces.

To enable MC-LAG, include the **mc-ae** statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level along with either the **ethernet-bridge**, **encapsulation ethernet-ccc**, **encapsulation ethernet-vpls**, or **flexible-ethernet-services** statement at the **[edit interfaces ae*X*]** hierarchy level. You also need to configure the **lacp** statement and the **admin-key** and **system-id** statements at the **[edit interfaces ae*X* aggregated-ether-options] hierarchy level**:

```
[edit interfaces aeX]
encapsulation (ethernet-bridge | ethernet-ccc | ethernet-vpls | flexible-ethernet-services);
aggregated-ether-options {
  lacp {
    active;
    admin-key number;
    system-id mac-address;
    system-priority number;
  }
  mc-ae {
    chassis-id chassis-id;
    events {
      iccp-peer-down {
        force-icl-down;
        prefer-status-control-active;
      }
    }
```

```
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
    }
```

NOTE: When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning.

To delete a MC-LAG interface from the configuration, issue the **delete interfaces ae*X* aggregated-ether-options mc-ae** command at the **[edit]** hierarchy level in configuration mode:

```
[edit]
user@host# delete interfaces aeX aggregated-ether-options mc-ae
```

Preventing Loops in MC-LAG Topologies

To prevent loops in MC-LAG topologies, configure the two edge nodes with same (STP) virtual root ID using Reverse L2 Gateway Protocol (RL2GP). This root ID should be superior to all bridges in the downstream network while downstream bridges have to be capable of running STP. RL2GP should be configured on both MC-LAG nodes to prevent loops. A potential loop, such as one that can happen due to improper cabling at the core or the access switching layer, or due to a bug in server software, is broken by STP blocking one of the interfaces in the downstream network. Because both the MC-LAG nodes are root bridges (virtual), the MC-LAG interface remains in the forwarding state. Downstream bridge receives bridge protocol data units (BPDUs) from both the nodes and thus receives twice the number of BPDUs on its **ae** interface. If both MC-LAG nodes use the same **ae** interface name, STP port number will be identical, which reduces the STP load on downstream bridge.

Configuring MC-LAG Devices

Perform the following steps on each switch that is hosting an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-options mc-ae mc-ae-id mc-ae-id
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
   ```
2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-options mc-ae chassis-id chassis-id
   ```


For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.

NOTE: Only active-active mode is supported for Reverse Layer 2 Gateway Protocol (R-L2GP) at this time.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae mode mode
```

For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options mc-ae mode active-active
```

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary. Primary is **active**, and secondary is **standby**.

NOTE: You must configure status control on both switches hosting the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
user@host# set aeX aggregated-ether-options mc-ae status-control active
```

5. Specify the same LACP system ID on each switch.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options lacp system-id mac-address
```

For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

6. Specify the same LACP administration key on each switch.

```
[edit interfaces]
user@host# set aeX aggregated-ether-options lacp admin-key number
```

For example:

```
[edit interfaces]
user@host# set ae1 aggregated-ether-options lacp admin-key 3
```

7. Configure ICCP by doing the following on each switch hosting the MC-LAG:

   a. Configure the local IP address to be used by all switches hosting the MC-LAG.

      [edit protocols]
      user@host# **set iccp local-ip-addr *local-ip-address***

      For example:

      [edit protocols]
      user@host# **set iccp local-ip-addr 3.3.3.1**

   b. (Optional) Configure the IP address of the router and the time during which an ICCP connection must succeed between the routeres hosting the MC-LAG.

      Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

      [edit protocols]
      user@host# **set iccp peer *peer-ip-address* session-establishment-hold-time *seconds***

      For example:

      [edit protocols]
      user@host# **set iccp peer 3.3.3.2 session-establishment-hold-time 50**

   c. (Optional) Configure the IP address to be used for backup liveness detection:

      NOTE: By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

      [edit protocols]
      user@host# **set iccp peer *peer-ip-address* backup-liveness-detection backup-peer-ip *ip-address***

      For example:

      [edit protocols]
      user@host# **set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232**

   d. Configure the minimum interval at which the router must receive a reply from the other router with which it has established a Bidirectional Forwarding Detection (BFD) session.

      NOTE: Configuring the minimum receive interval is required to enable BFD.

      [edit protocols]
      user@host# **set iccp peer *peer-ip-address* liveness-detection minimum-receive-interval *seconds***

      For example:

      [edit protocols]
      user@host# **set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60**

   e. Configure the minimum transmit interval during which a router must receive a reply from a router with which it has established a BFD session.


```
[edit protocols]
user@host# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval seconds
```

For example:

```
[edit protocols]
user@host# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval
60
```

8. Configure a multichassis protection link between the routers.

```
[edit]
user@host# set multi-chassis multi-chassis-protection peer-ip-address interface
interface-name
```

For example:

```
[edit protocols]
user@host# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

9. Enable RSTP globally on all interfaces.

```
[edit]
user@host# set protocols rstp interface all mode point-to-point
```

10. Disable RSTP on the ICL-PL interfaces on both routeres.

```
[edit]
user@host# set protocols rstp interface interface-name disable
```

For example:

```
[edit]
user@host# set protocols rstp interface ae0.0 disable
```

11. Configure the MC-LAG interfaces as edge ports on both routeres.

set protocols rstp interface *interface-name*

For example:

```
[edit]
user@host# set protocols rstp interface ae1.0
```

12. Enable BPDU block on all interfaces except for the ICL-PL interfaces on both routeres.

```
[edit]
user@host# set protocols rstp bpdu-block-on-edge
```

For example:

```
[edit]
user@host# set protocols rstp bpdu-block-on-edge
```

Understanding the Incremented Values of Statistical Counters for Loop-Free MC-LAG Networks

In an MC-LAG in an active-active bridging domain, the output of the following command displays the MC-LAG color counters to be continuously increasing. This increase in the statistical count is an expected behavior because the MC-LAG color attribute or counter functions as a loop prevention mechanism.

```
request pfe execute target fpc0 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   554712463 144488623417
request pfe execute target fpc0 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   554712747 144488664296
```

The exception table stored in the Packet Forwarding Engine contains a list of counters as displayed in the following example output:

```
request pfe execute target fpc0 command "show jnh 0 exceptions"
SENT: Ukern command: show jnh 0 exceptions
GOT: Reason                             Type         Packets      Bytes
GOT: ==================================================================
GOT: Ucode Internal
GOT: ----------------------
GOT: mcast stack overflow               DISC(33)           0          0
GOT: sample stack error                 DISC(35)           0          0
GOT: undefined nexthop opcode           DISC(36)           0          0
GOT: internal ucode error               DISC(37)           0          0
GOT: invalid fabric hdr version         DISC(41)           0          0
GOT:
GOT: PFE State Invalid
GOT: ----------------------
GOT: sw error                           DISC(64)   803092438 59795128826
GOT: child ifl nonlocal to pfe          DISC(85)           0          0
GOT: invalid fabric token               DISC(75)         179      42346
GOT: unknown family                     DISC(73)           0          0
GOT: unknown vrf                        DISC(77)           0          0
GOT: iif down                           DISC(87)           0          0
GOT: unknown iif                        DISC( 1)
GOT: invalid stream                     DISC(72)           0          0
GOT: egress pfe unspecified             DISC(19)       10889    1536998
GOT: invalid L2 token                   DISC(86)          26       1224
GOT: mc lag color                       DISC(88)   554693648
144486028726<<<<<<<<<<<<<<<<<<<<<<<<
GOT: dest interface non-local to pfe    DISC(27)  10939253797 2078273071209
GOT: invalid inline-svcs state          DISC(90)           0          0
GOT: nh id out of range                 DISC(93)           0          0
GOT: invalid encap                      DISC(96)           0          0
GOT: replication attempt on empty irb   DISC(97)           0          0
GOT: invalid selector entry             DISC(98)           0          0
GOT:
GOT:
GOT: Packet Exceptions
GOT: ----------------------
GOT: bad ipv4 hdr checksum              DISC( 2)
GOT: non-IPv4 layer3 tunnel             DISC( 4)           0          0
GOT: GRE unsupported flags              DISC( 5)           0          0
GOT: tunnel pkt too short               DISC( 6)           0          0
GOT: tunnel hdr too long                DISC( 7)           0          0
GOT: bad IPv6 options pkt               DISC( 9)           0          0
GOT: bad IP hdr                         DISC(11)           0          0
GOT: bad IP pkt len                     DISC(12)           0          0
GOT: L4 len too short                   DISC(13)
GOT: invalid TCP fragment               DISC(14)           0          0
GOT: mtu exceeded                       DISC(21)           0          0
GOT: frag needed but DF set             DISC(22)           0          0
GOT: ttl expired                        PUNT( 1)           9        769
GOT: IP options                         PUNT( 2)          16        512
GOT: xlated l2pt                        PUNT(14)           0          0
GOT: control pkt punt via ucode         PUNT( 4)           0          0
GOT: frame format error                 DISC( 0)
GOT: tunnel hdr needs reassembly        PUNT( 8)           0          0
GOT: GRE key mismatch                   DISC(76)           0          0
GOT: my-mac check failed                DISC(28)
GOT: frame relay type unsupported       DISC(38)           0          0
GOT: IGMP snooping control packet       PUNT(12)           0          0
```



```
GOT: bad CLNP hdr                      DISC(43)         0          0
GOT: bad CLNP hdr checksum             DISC(44)         0          0
GOT: Tunnel keepalives                 PUNT(58)         0          0
GOT:
GOT:
GOT: Bridging
GOT: ----------------------
GOT: lt unknown ucast                  DISC(84)         0          0
GOT: dmac miss                         DISC(15)         0          0
GOT: mac learn limit exceeded          DISC(17)         0          0
GOT: static mac on unexpected iif      DISC(18)         0          0
GOT: no local switching                DISC(20)         0          0
GOT: bridge ucast split horizon        DISC(26)     39458   13232394
GOT: mcast smac on bridged iif         DISC(24)      1263     200152
GOT: bridge pkt punt                   PUNT( 7)         0          0
GOT: iif STP blocked                   DISC( 3)
GOT: oif STP blocked                   DISC(31)
GOT: vlan id out of oif's range        DISC(32)
GOT: mlp pkt                           PUNT(11)  15188054  440453569
GOT: input trunk vlan lookup failed    DISC(91)         0          0
GOT: output trunk vlan lookup failed   DISC(92)         0          0
GOT: LSI/VT vlan validation failed     DISC(94)         0          0
GOT:
GOT:
GOT: Firewall
GOT: ----------------------
GOT: mac firewall                      DISC(78)
GOT: firewall discard                  DISC(67)         0          0
GOT: tcam miss                         DISC(16)         0          0
GOT: firewall reject                   PUNT(36)    155559   59137563
GOT: firewall send to host             PUNT(54)         0          0
GOT: firewall send to host for NAT     PUNT(59)         0          0
GOT:
GOT:
GOT: Routing
GOT: ----------------------
GOT: discard route                     DISC(66)   1577352   82845749
GOT: dsc ifl discard route             DISC(95)         0          0
GOT: hold route                        DISC(70)     21130    1073961
GOT: mcast rpf mismatch                DISC( 8)         0          0
GOT: resolve route                     PUNT(33)      2858     154202
GOT: control pkt punt via nh           PUNT(34)  51807272 5283911584
GOT: host route                        PUNT(32)  23473304 1370843994
GOT: ICMP redirect                     PUNT( 3)         0          0
GOT: mcast host copy                   PUNT( 6)         0          0
GOT: reject route                      PUNT(40)      1663     289278
GOT: link-layer-bcast-inet-check       DISC(99)         0          0
GOT:
```

Consider a sample deployment in which two provider edge (PE) devices, PE1 and PE2, are connected with an aggregated Ethernet (ae) interface, **ae0**. respectively. Multichassis link aggregation groups (MC-LAGs) are used between PE1 and PE2 to form a logical LAG interface between the two controllers. PE1 and PE2 in an MC-LAG use an interchassis control link-protection link (ICL-PL) to replicate forwarding information across the peers.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

The PE1 router is connected using another AE interface, **ae3**, to a host, H1, and to another MC-LAG host called C1. MC-LAG is enabled on the **ae3** interface.

Traffic received on PE1 from MC-LAG C1 can be flooded over the ICL to reach PE2. When the packets arrive at PE2, they can be flooded back to MC- LAG C1. Traffic sent by the single-homed host H1 can be flooded to MC-LAG C1 and the ICL on PE1. When PE2 receives such traffic from ICL, it can be again flooded to MC-LAG C1. To protect the MC-LAG topology from such loops, the MC-LAG color capability is implemented. This functionality is applied on the ingress of the ICL link. Therefore, when PE2 receives a packet from PE1, it sets the MC-LAG color as active or turns it on. When PE2 requires to flood the packet towards the MC-LAG link, it verifies whether the MC-LAG color bit is set or tagged as on. If the color is set, it drops the packet on the egress interface of MC-LAG ae3 member link interfaces and the **mc-lag color** counter in the jnh exceptions is incremented.

Such a behavior of increase in counter value is an expected condition in an MC-LAG configured in an active-active bridging domain and when any form of traffic that needs to be flooded, such as ARP broadcast or multicast traffic, traverses the network.

Every VLAN might drop some packets to prevent loops and such a drop of packets might not be specific to a VLAN.

Sometimes, on both MC LAGs of the MX Series routers, you might notice that the counter increases on FPC0 and FPC2, but it does not increase on FPC2 as illustrated in the following sample output:

```
request pfe execute target fpc0 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   558477875 144977739683
request pfe execute target fpc1 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)           0            0
request pfe execute target fpc2 command "show jnh 0 exceptions" |grep color
GOT: mc lag color                       DISC(88)   518499257 119130527834
```

This behavior occurs because four packet Forwarding Engines are present on MX Series routers with the 16-port 10-Gigabit Ethernet MPC (16x10GE 3D MPC) PFEs. If you examine one Packet Forwarding Engine in FPC 0, 1, and 2, PFE1 in FPC1 does not have any interfaces which are member of MC-LAG. It might contain interfaces in other aggregated Ethernet interfaces that are are not part of MC-LAG. Therefore, to obtain the correct counter statistics, you must examine the other PFEs by entering the **request pfe execute target fpc0 command "show jnh X exceptions" |grep color** command where X can be 0, 1, 2, or 3.

When the counter named **dest interface non-local to pfe** is increasing, it is a desired behavior when aggregated Ethernet interfaces are split over more than one FPC. Consider an example in which an **ae5** interface contains the following member links: **xe-0/1/0** on (FPC0) and **xe-1/1/0** (FPC1) Based on the hash algorithm, traffic must be split between these two links. The hash algorithm is applied on the ingress FPC and performs an operation where it marks each packet through which FPC must be forwarded (FPC0 or FPC1). Then the packet is sent to the fabric. From the fabric, all of traffic is sent to both FPCs 0 and 1. On FPC0, the microkernel analyzes the packet and determines whether the packet needs to be forwarded by the local interface (local to pfe) or whether this packet has already been forwarded through FPC1 (non-local to pfe). If the packet has been already forwarded, the packet is dropped and the **non-local to pfe**counter is incremented.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on page 127
- **show interfaces mc-ae on page 378**
- *Ethernet Interfaces*

Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers

The following sections describe the configuration of active-active bridging and VRRP over IRB in multichassis link aggregation (MC-LAG) on MX Series routers:

Configuring MC-LAG

An MC-LAG is composed of logical link aggregation groups (LAGs) and is configured under the **[edit interfaces ae*X*]** hierarchy, as follows:

```
[edit]
interfaces {
  ae0 {
    encapsulation ethernet-bridge;
    multi-chassis-protection {
      peer 10.10.10.10 {
        interface ge-0/0/0;
        }
      }
    aggregated-ether-options {
      mc-ae {
        mode active-active; # see note below
      }
    }
  }
```

NOTE: The **mode active-active** statement is valid only if **encapsulation** is **ethernet-bridge** or **extended-vlan-bridge**.

Use the **mode** statement to specify if a MC-LAG is **active-standby** or **active-active**. If the ICCP connection is UP and ICL comes UP, the router configured as standby will bring up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer.

Using **multi-chassis-protection** at the physical interface level is a way to reduce the configuration at the logical interface level.

If the following assumption exists (follow the above example):

If there are n+1 logical interfaces under **ae0**, from **ae0.0** through **ae0.n**, there will be n+1 logical interfaces under **ge-0/0/0** as well, from **ge-0/0/0.0** through **ge-0/0/0.n**, each **ge-0/0/0** logical interface will be a protection link for the **ae0** logical interface.

NOTE: **A bridge domain cannot have MC-AE logical interfaces which belong to different redundancy groups.**

Configuring Interchassis Link Label

The interchassis link-protection link (ICL-PL) provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL-PL is an aggregated Ethernet interface. You can configure only one ICL-PL between the two peers, although you can configure multiple MC-LAGs between them.

Ethernet as interchassis link label (ICL-PL) (assumes interface **ge-0/0/0.0** is used to protect interface **ae0.0** of MC-LAG-1):

```
[edit]
interfaces {
  ae0 {
    ....
    unit 0 {
      multi-chassis-protection {
        peer 10.10.10.10 {
          interface ge-0/0/0.0;
        }
        ....
      }
      ...
    }
  }
}
```

The protection interface can be an Ethernet type interface like **ge**, **xe**, or an aggregated Ethernet (**ae**) interface.

Configuring Multiple Chassis

A top-level hierarchy is used to specify multichassis-related configuration, as follows:

```
[edit]
multi-chassis {
  multi-chassis-protection {
    peer 10.10.10.10 {
```



```
            interface ge-0/0/0;
        }
    }
}
```

The above example specifies interface **ge-0/0/0** as the multichassis protection interface for all the multichassis aggregated Ethernet (MC-AE) interfaces which are also part of the peer. This can be overridden by specifying protection at the physical interface level and the logical interface level.

Configuring Service ID

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. You can configure the service IDs under the level of the hierarchies shown in the following examples:

Global Configuration (default logical system)

```
switch-options {
    service-id 10;
}
bridge-domains {
    bd0 {
        service-id 2;
    }
}
routing-instances {
    r1 {
        switch-options {
            service-id 10;
        }
        bridge-domains {
            bd0 {
                service-id 2;
            }
        }
    }
}
```

Logical Systems

```
logical-system {
    ls1 {
        switch-options {
            service-id 10;
        }
    }
}
logical-system {
    ls1 {
        routing-instances {
            r1 {
                switch-options {
                    service-id 10;
                }
            }
        }
    }
}
```

NOTE: Using a service name per bridge domain is not supported.

The bridge level service ID is required to link related bridge domains across peers, and should be configured with the same value. The **service-id** values share the name space across all bridging and routing instances, and across peers. Thus, duplicate values for service IDs are not permitted across these entities.

The service ID at the bridge domain level is mandatory for type non-single VLAN bridge domains. The service ID is optional for bridge domains with a VID defined. If no service ID is defined in the latter case, it is picked up from the service ID configuration for that routing instance.

NOTE: When this default routing instance (or any other routing instance) which contains a bridge domain containing an MC-AE interface is configured, you must configure a global level **switch-options service-id** *number*, irrespective of whether the contained bridge domains have specific service IDs configured.

In the example shown in Figure 12 on page 54, network routing instances N1 and N2, both for the same service ID, are configured with same service-id in both N1 and N2. Use of a name string instead of a number is not supported.

Figure 12: N1 and N2 for the Same Service with Same Service ID

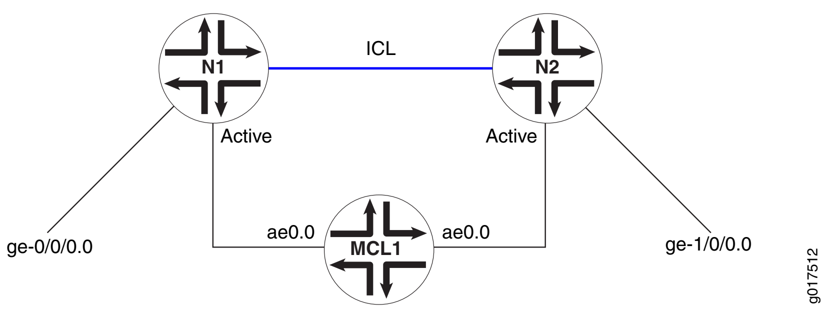

The following configuration restrictions apply:

- The service ID must be configured when the MC-AE interface is configured and an MC-AE logical interface is part of a bridge domain. This requirement is enforced.
- A single bridge domain cannot correspond to two redundancy group IDs.

Figure 13: Bridge Domain with Logical Interfaces from Two MC-AE Interfaces

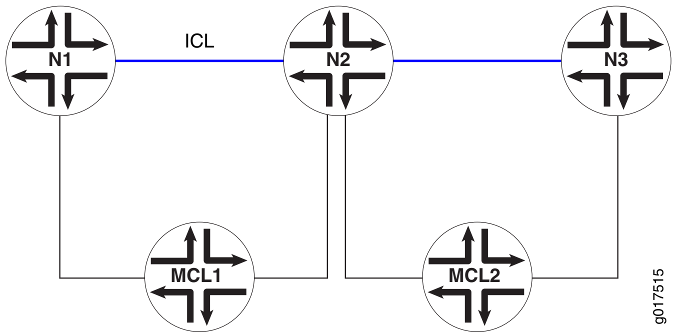

In Figure 13 on page 55, it is possible to configure a bridge domain consisting of logical interfaces from two MC-AE interfaces and map them to a separate redundancy group ID, which is not supported. A service should be mapped one-to-one with the redundancy group providing the service. This requirement is enforced.

To display the MC-AE configuration, use the **show interfaces *mc-ae*** command. For more information, see the CLI Explorer.

Configuring IGMP Snooping for Active-Active MC-LAG

For the multicast solution to work, the following must be configured:

- The multichassis protection link should be configured as a router-facing interface.

```
[edit bridge-domain bd-name]
protocols {
  igmp-snooping {
    interface ge-0/0/0.0 {
      multicast-router-interface;
    }
  }
}
```

  In this example, **ge-0/0/0.0** is an ICL interface.

- The **multichassis-lag-replicate-state** statement options should be configured under the **multicast-snooping-options** statement for that bridge domain.

NOTE: Snooping with active-active MC-LAG is only supported in non-proxy mode.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 56
- Configuring Multichassis Link Aggregation on page 42

- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on page 127
- **mc-ae on page 237**
- **multi-chassis-protection on page 242**
- **peer on page 244**
- **show interfaces irb on page 358**
- **show interfaces mc-ae on page 378**
- *Ethernet Interfaces*

Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers

This example illustrates how to configure a multichassis link aggregation group (MC-LAG) in an active-active scenario on MX Series routers.

- Requirements on page 56
- Overview on page 56
- Configuring the PE Routers on page 57
- Configuring the CE Router on page 67
- Configuring the Provider Router on page 69
- Verification on page 71

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Overview

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

In Figure 14 on page 57, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facindg the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

As a best practice, we recommend that you configure the ICCP and ICL interfaces over aggregated Ethernet interfaces instead of other interfaces such as Gigabit Ethernet interfaces, depending on your topology requirements and framework. You must disable RSTP on the ICL-PL interfaces for an MC-LAG in an active-active bridging domain.

Topology Diagram

Figure 14 on page 57 shows the topology used in this example. The interface **ge-1/0/2** functions as the ICCP link between the two PE devices, interface **ge-1/1/1** is the ICL-PL link, and interface **ge-1/1/4** is the link that connects to the server or the MC- LAG client device.

Figure 14: MC-LAG Active-Active on MX Series Routers

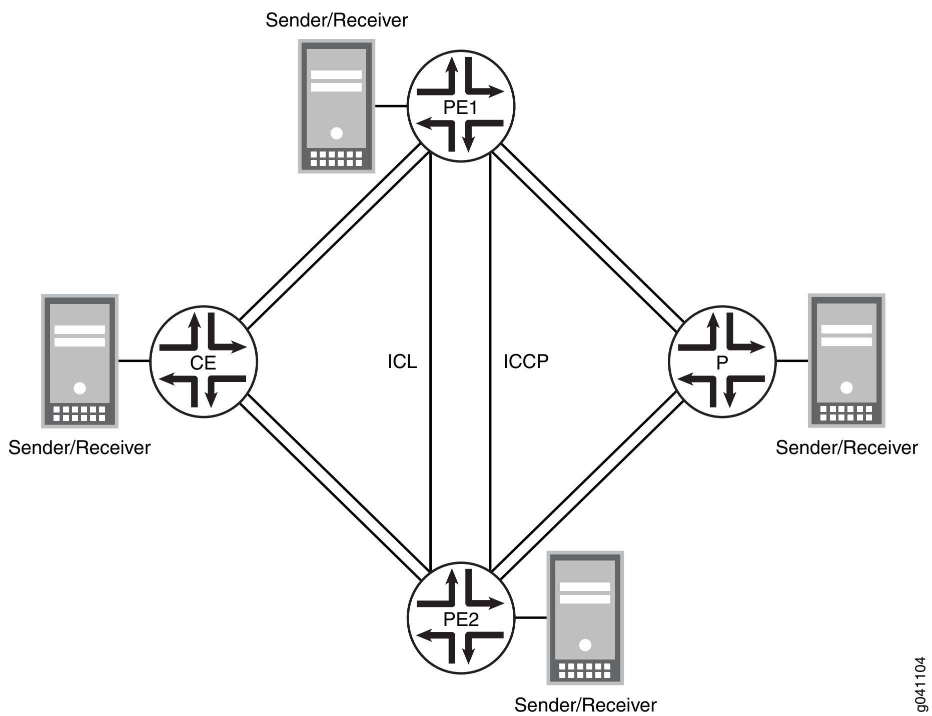

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.2
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

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

   ```
   [edit chassis]
   user@PE1# set aggregated-devices ethernet device-count 5
   ```

2. Specify the members to be included within the aggregated Ethernet bundles.

   ```
   [edit interfaces]
   user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
   user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
   ```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

   ```
   [edit interfaces]
   user@PE1# set ge-1/1/1 flexible-vlan-tagging
   user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
   user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
   user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
   user@PE1# set ge-1/1/4 flexible-vlan-tagging
   user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
   user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
   user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
   user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
   ```

4. Configure parameters on the aggregated Ethernet bundles.

   ```
   [edit interfaces ae0]
   user@PE1# set flexible-vlan-tagging
   user@PE1# set encapsulation flexible-ethernet-services
   user@PE1# set unit 0 encapsulation vlan-bridge
   user@PE1# set unit 0 vlan-id-range 100-110
   user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

   [edit interfaces ae1]
   user@PE1# set flexible-vlan-tagging
   user@PE1# set encapsulation flexible-ethernet-services
   user@PE1# set unit 0 encapsulation vlan-bridge
   user@PE1# set unit 0 vlan-id-range 100-110
   user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
   ```

5. Configure LACP on the aggregated Ethernet bundles.

   ```
   [edit interfaces ae0 aggregated-ether-options]
   user@PE1# set lacp active
   user@PE1# set lacp system-priority 100
   user@PE1# set lacp system-id 00:00:00:00:00:05
   user@PE1# set lacp admin-key 1
   ```

```
[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 59).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 59) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To configure Router PE2:

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

```
[edit chassis]
user@PE2# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae1
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE2# set ge-1/0/3 flexible-vlan-tagging
user@PE2# set ge-1/0/3 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/3 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/3 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/4 flexible-vlan-tagging
user@PE2# set ge-1/0/4 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/4 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/4 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae0
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae1
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```



```
[edit interfaces ae1]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1

[edit interfaces ae1 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 5
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 10
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE2 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@PE2# set domain-type bridge
   user@PE2# set vlan-id all
   user@PE2# set service-id 20
   user@PE2# set interface ae0.0
   user@PE2# set interface ae1.0
   user@PE2# set interface ge-1/0/3.0
   user@PE2# set interface ge-1/1/1.0
   user@PE2# set interface ge-1/1/4.0
   ```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

   The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

   ```
   [edit protocols iccp]
   user@PE2# set local-ip-addr 100.100.100.2
   user@PE2# set peer 100.100.100.1 redundancy-group-id-list 10
   user@PE2# set peer 100.100.100.1 liveness-detection minimum-interval 1000
   ```

9. Configure the service ID at the global level.

   ```
   [edit switch-options]
   user@PE2# set service-id 10
   ```

   You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
}
```

```
user@PE1# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
```



```
    }
}
user@PE1# show interfaces
ge-1/0/1 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/0/2 {
    unit 0 {
        family inet {
            address 100.100.100.1/30;
        }
    }
}
ge-1/1/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
ge-1/1/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 5;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
```



```
            encapsulation vlan-bridge;
            vlan-id-range 100-110;
            multi-chassis-protection 100.100.100.2 {
                interface ge-1/1/4.0;
            }
        }
    }
    ae1 {
        flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        aggregated-ether-options {
            lacp {
                active;
                system-priority 100;
                system-id 00:00:00:00:00:05;
                admin-key 1;
            }
            mc-ae {
                mc-ae-id 10;
                redundancy-group 10;
                chassis-id 1;
                mode active-active;
                status-control active;
            }
        }
        unit 0 {
            encapsulation vlan-bridge;
            vlan-id-range 100-110;
            multi-chassis-protection 100.100.100.2 {
                interface ge-1/1/4.0;
            }
        }
    }

    user@PE1# show protocols
    iccp {
        local-ip-addr 100.100.100.1;
        peer 100.100.100.2 {
            redundancy-group-id-list 10;
            liveness-detection {
                minimum-interval 1000;
            }
        }
    }

    user@PE1# show switch-options
    service-id 10;
```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

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

   ```
   [edit chassis]
   user@CE# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
   user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-2/1/6]
   user@CE# set flexible-vlan-tagging
   user@CE# set encapsulation flexible-ethernet-services
   user@CE# set unit 0 encapsulation vlan-bridge
   user@CE# set unit 0 vlan-id-range 100-110
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae0]
   user@CE# set flexible-vlan-tagging
   ```

```
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}
```

```
user@CE# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
  gigether-options {
    802.3ad ae0;
  }
}
```



```
ge-2/0/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/1/6 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

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

   ```
   [edit chassis]
   user@P# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@P# set ge-1/0/5 gigether-options 802.3ad ae1
   user@P# set ge-1/0/11 gigether-options 802.3ad ae1
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-1/1/3]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-500
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-110
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1 aggregated-ether-options]
   user@P# set lacp active
   user@P# set lacp system-priority 100
   ```

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@P# set vlan-id all
   user@P# set domain-type bridge
   user@P# set interface ge-1/1/3.0
   user@P# set interface ae1.0
   ```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
```



```
    domain-type bridge;
    vlan-id all;
    interface ge-1/1/3.0;
    interface ae1.0;
}
user@P# show chassis
aggregated-devices {
    ethernet {
        device-count 2;
    }
}
user@P# show interfaces
ge-1/0/5 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/11 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/1/3 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-500;
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**

- **show interfaces ae1**
- **show interfaces mc-ae**
- **show l2-learning instance extensive**

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- *Configuring ICCP for MC-LAG*
- **show interfaces (Aggregated Ethernet) on page 265** in the CLI Explorer

Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on Logical Systems on MX Series Routers

This example illustrates how to configure a multichassis link aggregation group (MC-LAG) in an active-active scenario on logical systems within MX Series routers.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Overview

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. CE is configured with a logical system, LS3. PE1 is configured with a logical system, LS1, and PE2 is configured with a logical system, LS2. LS1 and LS2 on the two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In this example, the logical system on the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. LS1 and LS2 on the two PE devices each have a LAG connected to LS3 on the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

In Figure 14 on page 57, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and LS3 on the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facing the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two logical systems, LS1 and LS2, of the PE devices. In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

As a best practice, we recommend that you configure the ICCP and ICL interfaces over aggregated Ethernet interfaces instead of other interfaces such as Gigabit Ethernet interfaces, depending on your topology requirements and framework. You must disable RSTP on the ICL-PL interfaces for an MC-LAG in an active-active bridging domain.

Topology Diagram

Figure 14 on page 57 shows the topology used in this example. The interface **ge-1/0/2** functions as the ICCP link between the two PE devices, interface **ge-1/1/1** is the ICL-PL link, and interface **ge-1/1/4** is the link that connects to the server or the MC- LAG client device. In the following figure, the logical systems, LS1 and LS2, on the PE devices, PE1 and PE2, are not explicitly displayed. Similarly, LS3 on the CE device is not shown in the illustration. However, the three logical systems are configured on the corresponding devices.

Figure 15: MC-LAG Active-Active on MX Series Routers

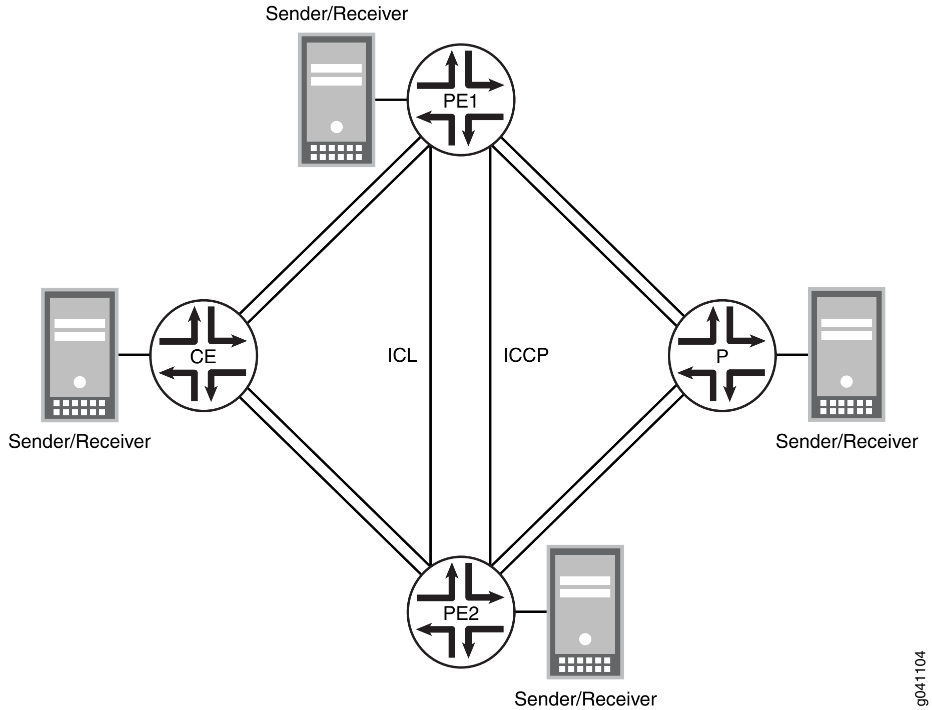

Configuring the Logical Systems on the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

LS1 on Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set logical-systems LS1 interfaces ge-1/0/1 gigether-options 802.3ad ae1
set logical-systems LS1 interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set logical-systems LS1 interfaces ge-1/0/6 gigether-options 802.3ad ae0
set logical-systems LS1 interfaces ge-1/1/1 flexible-vlan-tagging
set logical-systems LS1 interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set logical-systems LS1 interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set logical-systems LS1 interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set logical-systems LS1 interfaces ge-1/1/4 flexible-vlan-tagging
set logical-systems LS1 interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set logical-systems LS1 interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set logical-systems LS1 interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set logical-systems LS1 interfaces ae0 flexible-vlan-tagging
set logical-systems LS1 interfaces ae0 encapsulation flexible-ethernet-services
set logical-systems LS1 interfaces ae0 aggregated-ether-options lacp active
set logical-systems LS1 interfaces ae0 aggregated-ether-options lacp system-priority
  100
set logical-systems LS1 interfaces ae0 aggregated-ether-options lacp system-id
  00:00:00:00:00:05
set logical-systems LS1 interfaces ae0 aggregated-ether-options lacp admin-key 1
set logical-systems LS1 interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
```



```
set logical-systems LS1 interfaces ae0 aggregated-ether-options mc-ae redundancy-group
  10
set logical-systems LS1 interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set logical-systems LS1 interfaces ae0 aggregated-ether-options mc-ae mode
  active-active
set logical-systems LS1 interfaces ae0 aggregated-ether-options mc-ae status-control
  active
set logical-systems LS1 interfaces ae0 unit 0 encapsulation vlan-bridge
set logical-systems LS1 interfaces ae0 unit 0 vlan-id-range 100-110
set logical-systems LS1 interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2
  interface ge-1/1/4.0
set logical-systems LS1 interfaces ae1 flexible-vlan-tagging
set logical-systems LS1 interfaces ae1 encapsulation flexible-ethernet-services
set logical-systems LS1 interfaces ae1 aggregated-ether-options lacp active
set logical-systems LS1 interfaces ae1 aggregated-ether-options lacp system-priority 100
set logical-systems LS1 interfaces ae1 aggregated-ether-options lacp system-id
  00:00:00:00:00:05
set logical-systems LS1 interfaces ae1 aggregated-ether-options lacp admin-key 1
set logical-systems LS1 interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set logical-systems LS1 interfaces ae1 aggregated-ether-options mc-ae redundancy-group
  10
set logical-systems LS1 interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set logical-systems LS1 interfaces ae1 aggregated-ether-options mc-ae mode active-active
set logical-systems LS1 interfaces ae1 aggregated-ether-options mc-ae status-control
  active
set logical-systems LS1 interfaces ae1 unit 0 encapsulation vlan-bridge
set logical-systems LS1 interfaces ae1 unit 0 vlan-id-range 100-110
set logical-systems LS1 interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2
  interface ge-1/1/4.0
set logical-systems LS1 bridge-domains bd0 domain-type bridge
set logical-systems LS1 bridge-domains bd0 vlan-id all
set logical-systems LS1 bridge-domains bd0 service-id 20
set logical-systems LS1 bridge-domains bd0 interface ae1.0
set logical-systems LS1 bridge-domains bd0 interface ge-1/0/3.0
set logical-systems LS1 bridge-domains bd0 interface ge-1/1/1.0
set logical-systems LS1 bridge-domains bd0 interface ge-1/1/4.0
set logical-systems LS1 bridge-domains bd0 interface ae0.0
set logical-systems LS1 protocols iccp local-ip-addr 100.100.100.1
set logical-systems LS1 protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set logical-systems LS1 protocols iccp peer 100.100.100.2 liveness-detection
  minimum-interval 1000
set logical-systems LS1 switch-options service-id 10
```

LS2 on Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set logical-systems LS2 interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set logical-systems LS2 interfaces ge-1/0/3 flexible-vlan-tagging
set logical-systems LS2 interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set logical-systems LS2 interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set logical-systems LS2 interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set logical-systems LS2 interfaces ge-1/0/4 flexible-vlan-tagging
set logical-systems LS2 interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set logical-systems LS2 interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set logical-systems LS2 interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set logical-systems LS2 interfaces ge-1/0/5 gigether-options 802.3ad ae0
set logical-systems LS2 interfaces ge-1/1/0 gigether-options 802.3ad ae1
set logical-systems LS2 interfaces ae0 flexible-vlan-tagging
```

```
set logical-systems LS2 interfaces ae0 encapsulation flexible-ethernet-services
set logical-systems LS2 interfaces ae0 aggregated-ether-options lacp active
set logical-systems LS2 interfaces ae0 aggregated-ether-options lacp system-priority
  100
set logical-systems LS2 interfaces ae0 aggregated-ether-options lacp system-id
  00:00:00:00:00:05
set logical-systems LS2 interfaces ae0 aggregated-ether-options lacp admin-key 1
set logical-systems LS2 interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set logical-systems LS2 interfaces ae0 aggregated-ether-options mc-ae redundancy-group
  10
set logical-systems LS2 interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set logical-systems LS2 interfaces ae0 aggregated-ether-options mc-ae mode
  active-active
set logical-systems LS2 interfaces ae0 aggregated-ether-options mc-ae status-control
  active
set logical-systems LS2 interfaces ae0 unit 0 encapsulation vlan-bridge
set logical-systems LS2 interfaces ae0 unit 0 vlan-id-range 100-110
set logical-systems LS2 interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1
  interface ge-1/0/4.0
set logical-systems LS2 interfaces ae1 flexible-vlan-tagging
set logical-systems LS2 interfaces ae1 encapsulation flexible-ethernet-services
set logical-systems LS2 interfaces ae1 aggregated-ether-options lacp active
set logical-systems LS2 interfaces ae1 aggregated-ether-options lacp system-priority
  100
set logical-systems LS2 interfaces ae1 aggregated-ether-options lacp system-id
  00:00:00:00:00:05
set logical-systems LS2 interfaces ae1 aggregated-ether-options lacp admin-key 1
set logical-systems LS2 interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set logical-systems LS2 interfaces ae1 aggregated-ether-options mc-ae redundancy-group
  10
set logical-systems LS2 interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set logical-systems LS2 interfaces ae1 aggregated-ether-options mc-ae mode active-active
set logical-systems LS2 interfaces ae1 aggregated-ether-options mc-ae status-control
  active
set logical-systems LS2 interfaces ae1 unit 0 encapsulation vlan-bridge
set logical-systems LS2 interfaces ae1 unit 0 vlan-id-range 100-110
set logical-systems LS2 interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1
  interface ge-1/0/4.0
set logical-systems LS2 bridge-domains bd0 domain-type bridge
set logical-systems LS2 bridge-domains bd0 vlan-id all
set logical-systems LS2 bridge-domains bd0 service-id 20
set logical-systems LS2 bridge-domains bd0 interface ae1.0
set logical-systems LS2 bridge-domains bd0 interface ge-1/0/3.0
set logical-systems LS2 bridge-domains bd0 interface ge-1/0/4.0
set logical-systems LS2 bridge-domains bd0 interface ae0.0
set logical-systems LS2 protocols iccp local-ip-addr 100.100.100.2
set logical-systems LS2 protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set logical-systems LS2 protocols iccp peer 100.100.100.1 liveness-detection
  minimum-interval 1000
set logical-systems LS2 switch-options service-id 10
```


LS1 on Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the logical system, LS1, on router PE1:

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

```
[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit logical-systems LS1 interfaces]
user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit logical-systems LS1 interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit logical-systems LS1 interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

[edit logical-systems LS1 interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit logical-systems LS1 interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```



```
[edit logical-systems LS1 interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit logical-systems LS1 interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active

[edit logical-systems LS1 interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on LS1 of Router PE1 and LS2 of Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on LS1 of Router PE1 and LS2 of Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on LS2 of Router PE2, see "Router PE2" on page 59).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on LS1 of Router PE1 and LS2 of Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. LS1 of Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. LS2 of Router PE2 (as shown in "Router PE2" on page 59) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit logical-systems LS1 bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit logical-systems LS1 protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit logical-systems LS1 switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To configure LS2 of Router PE2:

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

```
[edit chassis]
user@PE2# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit logical-systems LS2 interfaces]
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae1
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit logical-systems LS2 interfaces]
user@PE2# set ge-1/0/3 flexible-vlan-tagging
user@PE2# set ge-1/0/3 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/3 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/3 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/4 flexible-vlan-tagging
user@PE2# set ge-1/0/4 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/4 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/4 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae0
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae1
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit logical-systems LS2 interfaces ae0]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
```



```
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0

[edit logical-systems LS2 interfaces ae1]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit logical-systems LS2 interfaces ae0 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1

[edit logical-systems LS2 interfaces ae1 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit logical-systems LS2 interfaces ae0 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 5
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active

[edit logical-systems LS2 interfaces ae1 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 10
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on LS1 of Router PE1 and LS2 of Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on LS1 of Router PE1 and LS2 of Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on LS1 of Router PE1 and LS2 of Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. LS1 of Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. LS2 of Router PE2 uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit logical-systems LS2 bridge-domains bd0]
user@PE2# set domain-type bridge
user@PE2# set vlan-id all
user@PE2# set service-id 20
user@PE2# set interface ae0.0
user@PE2# set interface ae1.0
user@PE2# set interface ge-1/0/3.0
user@PE2# set interface ge-1/1/1.0
user@PE2# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case LS1 of Router PE1 and LS2 of Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit logical-systems LS2 protocols iccp]
user@PE2# set local-ip-addr 100.100.100.2
user@PE2# set peer 100.100.100.1 redundancy-group-id-list 10
user@PE2# set peer 100.100.100.1 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit logical-systems LS2 switch-options]
user@PE2# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, , **show interfaces**, **show protocols**, and **show switch-options** commands at the **[edit logical-systems LS1]** hierarchy level, and the **show chassis** command at the **[edit]** hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit logical-systems LS1]
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
```

```
}
[edit]
user@PE1# show chassis
aggregated-devices {
   ethernet {
      device-count 5;
   }
}
[edit logical-systems LS1]
user@PE1# show interfaces
ge-1/0/1 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/6 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-1/0/2 {
   unit 0 {
      family inet {
         address 100.100.100.1/30;
      }
   }
}
ge-1/1/1 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ge-1/1/4 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
         system-id 00:00:00:00:00:05;
         admin-key 1;
      }
      mc-ae {
         mc-ae-id 5;
```

```
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}

[edit logical-systems LS1]
user@PE1# show protocols
iccp {
    local-ip-addr 100.100.100.1;
    peer 100.100.100.2 {
        redundancy-group-id-list 10;
        liveness-detection {
            minimum-interval 1000;
        }
    }
}

[edit logical-systems LS1]
user@PE1# show switch-options
service-id 10;
```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for LS2 of Router PE2, using the appropriate interface names and addresses.

Configuring the Logical System on the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

LS3 of Router CE

```
set chassis aggregated-devices ethernet device-count 2
set logical-systems LS3 interfaces ge-2/0/2 gigether-options 802.3ad ae0
set logical-systems LS3 interfaces ge-2/0/3 gigether-options 802.3ad ae0
set logical-systems LS3 interfaces ge-2/1/6 flexible-vlan-tagging
set logical-systems LS3 interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set logical-systems LS3 interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set logical-systems LS3 interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set logical-systems LS3 interfaces ae0 flexible-vlan-tagging
set logical-systems LS3 interfaces ae0 encapsulation flexible-ethernet-services
set logical-systems LS3 interfaces ae0 aggregated-ether-options lacp active
set logical-systems LS3 interfaces ae0 aggregated-ether-options lacp system-priority
  100
set logical-systems LS3 interfaces ae0 unit 0 encapsulation vlan-bridge
set logical-systems LS3 interfaces ae0 unit 0 vlan-id-range 100-500
set logical-systems LS3 bridge-domains bd0 domain-type bridge
set logical-systems LS3 bridge-domains bd0 vlan-id all
set logical-systems LS3 bridge-domains bd0 interface ge-2/1/6.0
set logical-systems LS3 bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

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

   ```
   [edit chassis]
   user@CE# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit logical-systems LS3 interfaces]
   user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
   user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit logical-systems LS3 interfaces ge-2/1/6]
   user@CE# set flexible-vlan-tagging
   user@CE# set encapsulation flexible-ethernet-services
   user@CE# set unit 0 encapsulation vlan-bridge
   ```

```
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit logical-systems LS3 interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit logical-systems LS3 interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

The **active** statement initiates transmission of LACP packets.

For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit logical-systems LS3 bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains** and **show interfaces** commands at the **[edit logical-systems LS3]** hierarchy level, and the **show chassis** command at the **[edit]** hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit logical-systems LS3]
user@CE# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-2/1/6.0;
   interface ae0.0;
}

[edit]
user@CE# show chassis
aggregated-devices {
   ethernet {
      device-count 2;
```

```
  }
}
[edit logical-systems LS3]
user@CE# show interfaces
ge-2/0/2 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/1/6 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
```

```
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

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

```
[edit chassis]
user@P# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@P# set ge-1/0/5 gigether-options 802.3ad ae1
user@P# set ge-1/0/11 gigether-options 802.3ad ae1
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-1/1/3]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-500
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae1]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-110
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae1 aggregated-ether-options]
user@P# set lacp active
user@P# set lacp system-priority 100
```

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@P# set vlan-id all
user@P# set domain-type bridge
user@P# set interface ge-1/1/3.0
user@P# set interface ae1.0
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   interface ge-1/1/3.0;
   interface ae1.0;
}
```

```
user@P# show chassis
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@P# show interfaces
ge-1/0/5 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/11 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/1/3 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
   }
}
ae1 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
      }
   }
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show l2-learning instance extensive**

Related Documentation

- Multichassis Link Aggregation on Logical Systems Overview on page 28
- *Configuring ICCP for MC-LAG*
- **show interfaces (Aggregated Ethernet) on page 265** in the CLI Explorer

Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the routers participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a router.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a router.

- Configure Virtual Router Redundancy Protocol (VRRP) on a router.

Overview

In this example, you configure an MC-LAG across two routers by including interfaces from both routers in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each router in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In Figure 14 on page 57, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facindg the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

Topology Diagram

Figure 14 on page 57 shows the topology used in this example.

Figure 16: MC-LAG Active-Active on MX Series Routers

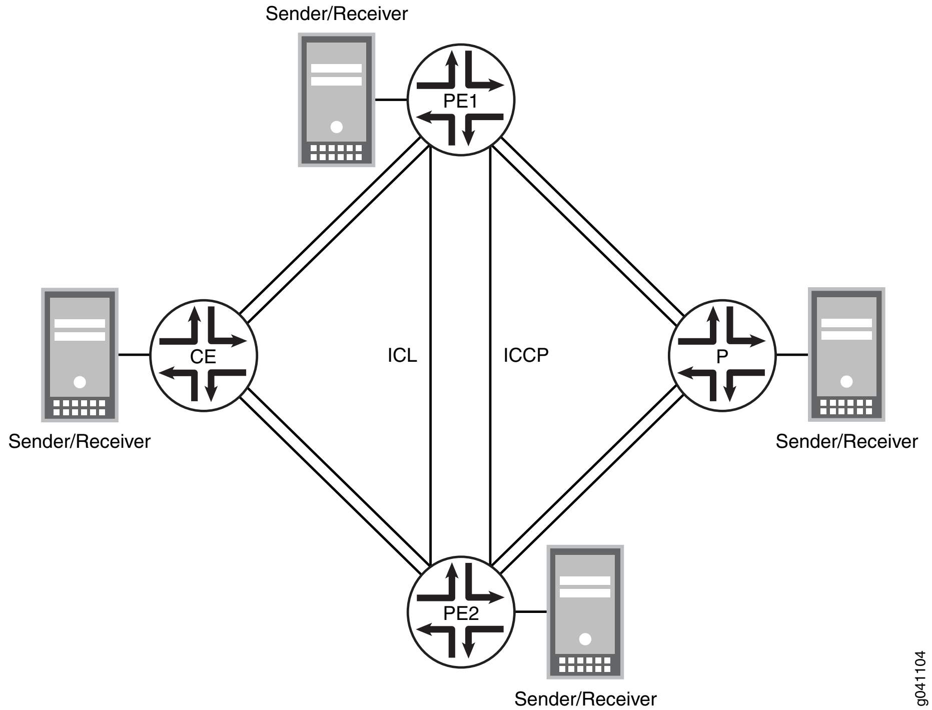

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface ge-1/1/4.0 priority 200
set protocols pim interface ge-1/1/4.0 version 2
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface ge-1/1/1.0 priority 600
set protocols pim interface ge-1/1/1.0 version 2
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
```

```
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.2
```



```
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface ge-1/0/4.0 priority 200
set protocols pim interface ge-1/0/4.0 version 2
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

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

```
[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 59).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 59) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs on PE1 and PE2:

1. Create a routed VLAN interface (RVI) for each MC-LAG, assign a virtual IP address that is shared between each router in the VRRP groups, and assign an individual IP address for each router in the VRRP groups.

PE1

```
[edit interfaces]
user@PE1# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1
   virtual-address 10.1.1.1
```

```
user@PE1# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2
  virtual-address 10.1.1.2
```

PE2

```
[edit interfaces]
user@PE2# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1
  virtual-address 10.1.1.1
user@PE2# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2
  virtual-address 10.1.1.2
```

2. Assign the priority for each router in the VRRP groups:

   NOTE: The router configured with the highest priority is the master.

PE1

```
[edit interfaces]
user@PE1# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority 200
user@PE1 set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority 200
```

PE2

```
[edit interfaces]
user@sPE2# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority
  150
user@PE2# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority
  150
```

3. Enable the router to accept all packets destined for the virtual IP address if it is the master in a VRRP group:

PE1

```
[edit interfaces]
user@PE1# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 accept-data
user@PE1# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 accept-data
```

PE2

```
[edit interfaces]
user@PE2# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 accept-data
user@PE2# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2
  accept-data
```

Step-by-Step Procedure

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area on PE1 and PE2.

```
[edit protocols ospf]
user@host# set area 0.0.0.0
```

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area on PE1 and PE2.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface ge-1/1/1.0
user@host# set interface ge-1/4/1.0
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces on PE1 and PE2.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
   700
user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
   minimum-interval 350
user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
   threshold 500
user@host# set interface ge-1/4/1.0 bfd-liveness-detection
   minimum-receive-interval 700
user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
   minimum-interval 350
user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
   threshold 500
```

Step-by-Step Procedure

To configure PIM as the multicast protocol:

1. Configure a static rendezvous point (RP) address on PE1 and PE2.

```
[edit protocols pim]
user@host# set rp static address 1.0.0.3
```

2. Configure the address ranges of the multicast groups for which PE1 and PE2 can be a rendezvous point (RP).

```
[edit protocols pim rp static address 1.0.0.3]
user@host# set group-ranges 239.0.0.0/8
```

3. Enable PIM on the VLAN interfaces for the MC-LAGs on PE1 and PE2.

```
[edit protocols pim]
user@host# set interface ge-1/1/1.0 version 2
user@host# set interface ge-1/4/1.0 version 2
```

4. Configure each PIM interface's priority for being selected as the designated router (DR).

   An interface with a higher priority value has a higher probability of being selected as the DR.

   PE1

```
[edit protocols pim]
user@host# set interface ge-1/1/1.0 priority 200
user@host# set interface ge-1/4/1.0 priority 600
```

   PE2

```
[edit protocols pim]
user@host# set interface ge-1/1/1.0 priority 100
user@host# set interface ge-1/4/1.0 priority 500
```


5. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the PIM interfaces on PE1 and PE2.

   ```
   [edit protocols pim]
   user@host# set interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
     700
   user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
     minimum-interval 350
   user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
     threshold 500
   user@host# set interface ge-1/4/1.0 bfd-liveness-detection
     minimum-receive-interval 700
   user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
     minimum-interval 350
   user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
     threshold 500
   ```

Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on PE1 and PE2.

   ```
   [edit]
   user@host# set protocols rstp interface all mode point-to-point
   ```

2. Disable RSTP on the ICL-PL interfaces on PE1 and PE2:

   ```
   [edit]
   user@host# set protocols rstp interface ae0.0 disable
   ```

3. Configure the MC-LAG interfaces as edge ports on PE1 and PE2.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

   ```
   [edit]
   user@host# set protocols rstp interface ae1.0 edge
   ```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on PE1 and PE2.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

   ```
   [edit]
   user@host# set protocols rstp bpdu-block-on-edge
   ```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
}
```

```
user@PE1# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
  }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/0/2 {
  unit 0 {
    family inet {
      address 100.100.100.1/30;
    }
  }
}
ge-1/1/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ge-1/1/4 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
```



```
        vlan-id-range 100-110;
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 5;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
```



```
user@PE1#show vrrp
vlan {
  unit 100 {
    family inet {
      address 10.1.1.11/24 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.21/24 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 200;
          accept-data;
        }
      }
    }
  }
}
```

```
user@PE1# show protocols
iccp {
  local-ip-addr 100.100.100.1;
  peer 100.100.100.2 {
    redundancy-group-id-list 10;
    liveness-detection {
      minimum-interval 1000;
    }
    rstp {
      interface ae0.0 {
        disable;
      }
      interface ae1.0 {
        edge;
      }
      interface all {
        mode point-to-point;
      }
      bpdu-block-on-edge;
    }
  }
}
rstp {
  interface ae0.0 {
    disable;
  }
  interface ae1.0 {
    edge;
  }
  interface all {
```

```
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
ospf {
    area 0.0.0.0 {
        interface ge-1/1/1.0 {
            bfd-liveness-detection {
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
            }
        }
        interface ge-1/4/1.0 {
            bfd-liveness-detection {
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
            }
        }
    }
}
pim {
    rp {
        static {
            address 1.0.0.3 {
                group-ranges {
                    239.0.0.0/8;
                }
            }
        }
    }
    interface ge-1/1/1.0 {
        priority 200;
        version 2;
        bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
            minimum-receive-interval 700;
            transmit-interval {
                minimum-interval 350;
                threshold 500;
            }
        }
    }
    interface ge-1/4/1.0 {
        priority 600;
        version 2;
        bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
            minimum-receive-interval 700;
            transmit-interval {
                minimum-interval 350;
                threshold 500;
            }
```

```
        }
      }
    }
```

user@PE1# **show switch-options**
service-id 10;

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

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

   [edit chassis]
   user@CE# **set aggregated-devices ethernet device-count 2**

2. Specify the members to be included within the aggregated Ethernet bundle.

   [edit interfaces]
   user@CE# **set ge-2/0/2 gigether-options 802.3ad ae0**
   user@CE# **set ge-2/0/3 gigether-options 802.3ad ae0**

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-2/1/6]
   user@CE# set flexible-vlan-tagging
   user@CE# set encapsulation flexible-ethernet-services
   user@CE# set unit 0 encapsulation vlan-bridge
   user@CE# set unit 0 vlan-id-range 100-110
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae0]
   user@CE# set flexible-vlan-tagging
   user@CE# set encapsulation flexible-ethernet-services
   user@CE# set unit 0 encapsulation vlan-bridge
   user@CE# set unit 0 vlan-id-range 100-500
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae0 aggregated-ether-options]
   user@CE# set lacp active
   user@CE# set lacp system-priority 100
   ```

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@CE# set domain-type bridge
   user@CE# set vlan-id all
   user@CE# set interface ge-2/1/6.0
   user@CE# set interface ae0.0
   ```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}

user@CE# show chassis
```

```
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/1/6 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
```

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

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

   ```
   [edit chassis]
   user@P# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@P# set ge-1/0/5 gigether-options 802.3ad ae1
   user@P# set ge-1/0/11 gigether-options 802.3ad ae1
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-1/1/3]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-500
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-110
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1 aggregated-ether-options]
   user@P# set lacp active
   user@P# set lacp system-priority 100
   ```

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@P# set vlan-id all
   user@P# set domain-type bridge
   user@P# set interface ge-1/1/3.0
   ```

```
user@P# set interface ae1.0
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-1/1/3.0;
  interface ae1.0;
}
```

```
user@P# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@P# show interfaces
ge-1/0/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/11 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/1/3 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show pim interfaces**
- **show vrrp**
- **show l2-learning instance extensive**

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 56
- Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers on page 109

Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router

Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the routers participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 unicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a router.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a router.
- Configure Virtual Router Redundancy Protocol (VRRP) on a router.

Overview

In this example, you configure an MC-LAG across two routers by including interfaces from both routers in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each router in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

In Figure 14 on page 57, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facindg the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

Topology Diagram

Figure 14 on page 57 shows the topology used in this example.

Figure 17: MC-LAG Active-Active on MX Series Routers

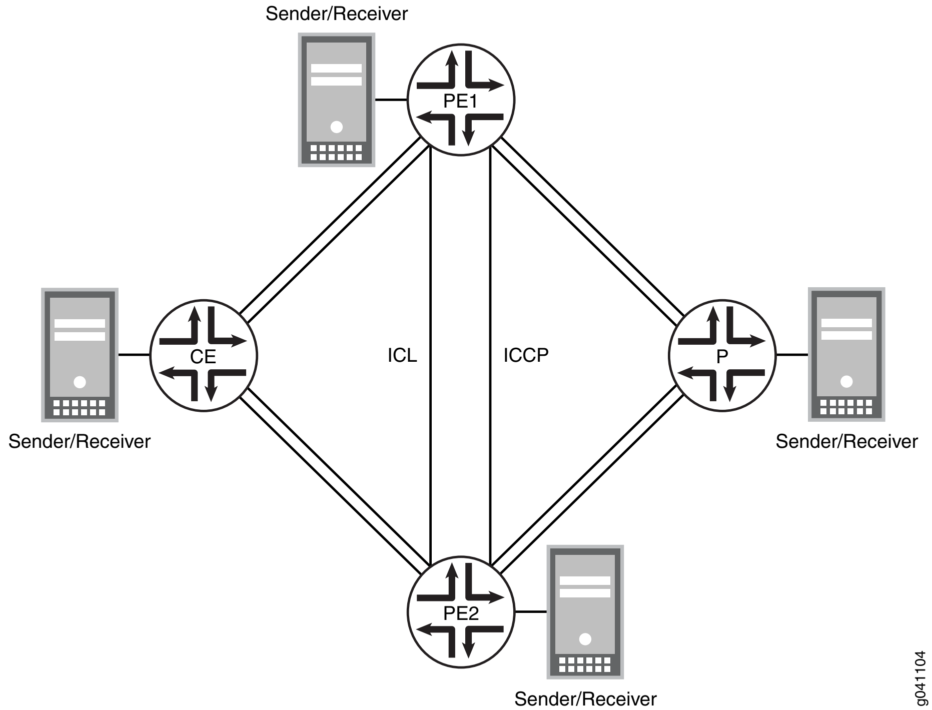

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```



```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
```

```
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.2
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

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

   ```
   [edit chassis]
   user@PE1# set aggregated-devices ethernet device-count 5
   ```

2. Specify the members to be included within the aggregated Ethernet bundles.

   ```
   [edit interfaces]
   user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
   user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
   ```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
```

```
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 59).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 59) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```


You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs on PE1 and PE2:

1. Enable VRRP on the MC-LAG on PE1 and PE2:

   - Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each router in the VRRP group, and assign an individual IP address for each router in the VRRP group:

     PE1:

     ```
     [edit interfaces]
     user@PE1# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 virtual-address 100.1.1.1
     ```

     PE2:

     ```
     [edit interfaces]
     user@PE2# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 virtual-address 100.1.1.1
     ```

   - Assign the priority for each router in the VRRP group:

     NOTE: The router configured with the highest priority is the master.

     PE1:

     ```
     [edit interfaces]
     user@PE1# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 priority 200
     ```

     PE2:

     ```
     [edit interfaces]
     user@PE2# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
     ```

   - Enable the router to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

     PE1:

     ```
     [edit interfaces]
     user@PE1# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 accept-data
     ```

     PE2:

     ```
     [edit interfaces]
     user@PE2# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
     ```

Step-by-Step Procedure

To enabled RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

   ```
   [edit]
   user@host# set protocols rstp interface all mode point-to-point
   ```

2. Disable RSTP on the ICL-PL interfaces on PE1 and PE2:


```
[edit]
user@host# set protocols rstp interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@host# set protocols rstp interface ae1.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@host# set protocols rstp bpdu-block-on-edge
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
}

user@PE1# show vrrp
vlan {
unit 100 {
  family inet {
    address 100.1.1.11/24 {
      vrrp-group 1 {
        virtual-address 100.1.1.1;
        priority 200;
        accept-data;
      }
    }
  }
}

user@PE1# show chassis
```

```
aggregated-devices {
   ethernet {
      device-count 5;
   }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/6 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-1/0/2 {
   unit 0 {
      family inet {
         address 100.100.100.1/30;
      }
   }
}
ge-1/1/1 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ge-1/1/4 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
         system-id 00:00:00:00:00:05;
         admin-key 1;
      }
      mc-ae {
         mc-ae-id 5;
         redundancy-group 10;
         chassis-id 1;
         mode active-active;
         status-control active;
```



```
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}

user@PE1# show protocols
iccp {
    local-ip-addr 100.100.100.1;
    peer 100.100.100.2 {
        redundancy-group-id-list 10;
        liveness-detection {
            minimum-interval 1000;
        }
        rstp {
            interface ae0.0 {
                disable;
            }
            interface ae1.0 {
                edge;
            }
            interface all {
                mode point-to-point;
            }
            bpdu-block-on-edge;
```

```
    }
  }
}
```

user@PE1# **show switch-options**
service-id 10;

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

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

   [edit chassis]
   user@CE# **set aggregated-devices ethernet device-count 2**

2. Specify the members to be included within the aggregated Ethernet bundle.

   [edit interfaces]
   user@CE# **set ge-2/0/2 gigether-options 802.3ad ae0**
   user@CE# **set ge-2/0/3 gigether-options 802.3ad ae0**

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-2/1/6]
   user@CE# set flexible-vlan-tagging
   user@CE# set encapsulation flexible-ethernet-services
   user@CE# set unit 0 encapsulation vlan-bridge
   user@CE# set unit 0 vlan-id-range 100-110
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae0]
   user@CE# set flexible-vlan-tagging
   user@CE# set encapsulation flexible-ethernet-services
   user@CE# set unit 0 encapsulation vlan-bridge
   user@CE# set unit 0 vlan-id-range 100-500
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae0 aggregated-ether-options]
   user@CE# set lacp active
   user@CE# set lacp system-priority 100
   ```

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@CE# set domain-type bridge
   user@CE# set vlan-id all
   user@CE# set interface ge-2/1/6.0
   user@CE# set interface ae0.0
   ```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}
```

```
user@CE# show chassis
```

```
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/1/6 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
```

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

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

   ```
   [edit chassis]
   user@P# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@P# set ge-1/0/5 gigether-options 802.3ad ae1
   user@P# set ge-1/0/11 gigether-options 802.3ad ae1
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-1/1/3]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-500
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-110
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1 aggregated-ether-options]
   user@P# set lacp active
   user@P# set lacp system-priority 100
   ```

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@P# set vlan-id all
   user@P# set domain-type bridge
   user@P# set interface ge-1/1/3.0
   ```

user@P# **set interface ae1.0**

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-1/1/3.0;
  interface ae1.0;
}
```

```
user@P# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@P# show interfaces
ge-1/0/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/11 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/1/3 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show vrrp**
- show l2-learning instance **extensive**

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 56
- Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers on page 89

Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces

In a multichassis link aggregation (MC-LAG) topology with active-standby mode, a link switchover happens only if the active node goes down. You can override this default behavior by configuring an MC-LAG interface in active-standby mode to automatically revert to a preferred node. With this configuration you can trigger a link switchover to a preferred node even when the active node is available. For example, consider two nodes, PE1 and PE2. PE1 is configured in active mode making it a preferred node and PE2 is configured in active-standby mode. In case of any failure at PE1, PE2 becomes the active node. However, as soon as PE1 is available again, an automatic link switchover is triggered and the control is switched back to PE1 even though PE2 is active.

You can configure the link switchover in two modes: revertive and nonrevertive. In revertive mode, the link switchover is triggered automatically while in nonrevertive mode the link switchover should be triggered manually by the user. You can also configure a revert time that triggers an automatic or manual switchover when the specified timer expires.

For nonrevertive mode, you can manually trigger a link switchover using the **request interface mc-ae switchover** operational mode command.

NOTE:

- If two MC-LAG devices configured in an active-standby setup using Interchassis Control Protocol (ICCP) and nonrevertive swithcover mode is configured on the aggregated Ethernet interfaces of both the MC-LAGs and when both mc-ae interfaces are linked together with Layer 2 circuit local-switching configuration, we recommend that you perform switchover by entering the **request interface mc-ae switchover (immediate mcae-id mcae-id | mcae-id mcae-id)** operational mode command only on one of the aggregated Ethernet interfaces of an MC-LAG device. This command can be issued only on MC-LAG devices that are configured as active nodes (by using the **status-control active** statement at the **[edit interfaces ae*X* aggregated-ether-options mc-ae]** hierarchy level).
- In nonrevertive switchover mode, when an MC-LAG interface transitions to the standby state because of an MC-LAG member link failure and another MC-LAG interface moves to the active state, the MC-LAG in standby state remains in that state until the MC-LAG in active state encounters a failure for it to return to the active state.
- If you perform a switchover on both the aggregated Ethernet interfaces in the MC-LAG, because of Layer 2 circuit local-switching configuration, a switchover on one aggregated Ethernet interface triggers a switchover on the other aggregated Ethernet interface. In such a scenario, both the aggregated Ethernet interfaces move to the standby state and then transition back to the active state. Therefore, you must not perform switchover on both the aggregated Ethernet interfaces in an MC-LAG at the same time.
- Layer 2 circuit configuration and VPLS functionalities are not supported if you configure an MC-LAG interface to be in revertive switchover mode. You can configure the revertive or nonrevertive switchover capability only if two MC-LAG devices are configured in an active-standby setup (one device set as an active node by using the status-control standby statement and the other device set as a standby node by using the **status-control active** statement at the **[edit interfaces ae*X* aggregated-ether-options mc-ae]** hierarchy level. You can perform a switchover by entering the **request interface mc-ae switchover (immediate mcae-id mcae-id | mcae-id mcae-id)** operational mode command only on MC-LAG devices configured as active nodes.

To configure the link switchover mechanism on an MC-LAG interface:

- Configure the link switchover in revertive mode.

  ```
  [edit interfaces aeX aggregated-ether-options mc-ae]
  user@host# set switchover-mode revertive;
  ```

- (Optional) Configure the link switchover in nonrevertive mode.

  ```
  [edit interfaces aeX aggregated-ether-options mc-ae]
  user@host# set switchover-mode non-revertive;
  ```


- Configure the revert time.

  [edit interfaces ae*X* aggregated-ether-options mc-ae]
  user@host# **set revert-time *revert-time***;

- Trigger manual switchover

  user@host> **request interface mc-ae switchover** {
  < **immediate**> mcae-id *mcae-id*;
  mcae-id *mcae-id*;
  }

You can use the **show interfaces mc-ae revertive-info** command to view the switchover configuration information.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Configuring Multichassis Link Aggregation on page 42
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 56

Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers

You can use the bridge-domain statement's service-id id option to specify the multichassis aggregated Ethernet configuration.

- The **service-id** statement is mandatory for non-single VLAN type bridge domains (**none**, **all** or **vlan-id-tags:dual**).
- It is optional for bridge domains with a VID defined.
- If no service-id is defined in the latter case, it will be picked up from the RTT's **service-id** configuration.
- The bridge level service-id is required to link related bridge domains across peers, and should be configured with the same value.
- The service-id values share the name space across all bridging and routing instances, and across peers. Thus, duplicate values for service-ids are not permitted across these entities.
- A change of bridge **service-id** is considered catastrophic, and the bridge domain is reincarnated.

This procedure allows you to enable or disable the replication feature. This option applies to all instances.

To configure IGMP snooping in active-standby mode:

1. Use the **multichassis-lag-replicate-state** statement at the **multicast-snooping-options** hierarchy level in the master instance.


```
multicast-snooping-options {
  ...
  multichassis-lag-replicate-state; # REQUIRED
}
```

The interchassis link, **interface *icl-intf-name***, of the learning domain should be a router facing interface.

1. Use the **interface *icl-intf-name*** statement at the **protocols igmp-snooping** hierarchy level, as shown in the following example:

```
protocols {
  igmp-snooping {
    interface icl-intf-name {
      multicast-router-interface;
    }
  }
}
```

Related Documentation

- IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview on page 20
- *Example: Configuring IGMP Snooping*
- **igmp-snooping on page 226**
- **multicast-router-interface on page 241**
- **show l2-learning instance on page 364**
- *Ethernet Interfaces*

Example: Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers

This example shows how to configure Internet Group Management Protocol (IGMP) snooping on MX Series routers with a multichassis link aggregation group (MC-LAG) in an active-active scenario.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Before you begin, make sure that Protocol Independent Multicast (PIM) and IGMP are running on all interfaces that will receive multicast packets. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM.

Overview

When links are aggregated, the links can be treated as if they were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. MC-LAG provides redundant Layer 2 access connectivity at the node level. This enables two or more systems to share a common LAG endpoint. The multiple end points present a single logical chassis to the start point, and the start node does not need to be aware that MC-LAG is being used.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

NOTE: **The other possible mode is active-standby, in which one of the router's ports only becomes active when failure is detected in the active links. In active-standby mode, the PE routers perform an election to determine the active and standby routers.**

In Figure 18 on page 133, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facing the CE device.

Internet Chassis Control (ICCP) control messages are sent between the two PE devices. These messages exchange MC-LAG configuration parameters and ensure that both chassis use the correct Link Aggregation Control Protocol (LACP) parameters when talking to the CE device.

The interchassis link (ICL) provides redundancy when a link failure occurs on one of the active links. The ICL-PL between the MC-LAG peering devices relays traffic that would otherwise be dropped due to a link failure.

Topology Diagram

Figure 18 on page 133 shows the topology used in this example.

Figure 18: IGMP Snooping in MC-LAG Active-Active on MX Series Routers

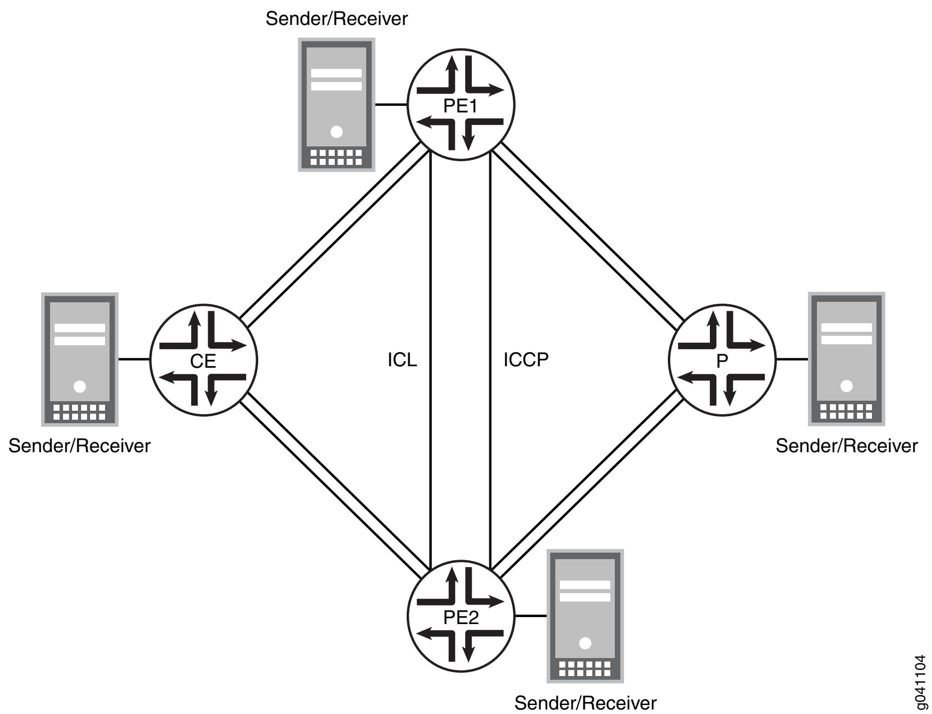

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set bridge-domains bd0 multicast-snooping-options multichassis-lag-replicate-state
set bridge-domains bd0 protocols igmp-snooping vlan 100 interface ge-1/1/4.0
  multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 101 interface ge-1/1/4.0
  multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 200 interface ge-1/1/4.0
  multicast-router-interface
set multicast-snooping-options multichassis-lag-replicate-state
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
```



```
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set bridge-domains bd0 multicast-snooping-options multichassis-lag-replicate-state
set bridge-domains bd0 protocols igmp-snooping vlan 100 interface ge-1/0/4.0
  multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 101 interface ge-1/0/4.0
  multicast-router-interface
set bridge-domains bd0 protocols igmp-snooping vlan 200 interface ge-1/0/4.0
  multicast-router-interface
set multicast-snooping-options multichassis-lag-replicate-state
set protocols iccp local-ip-addr 100.100.100.2
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

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

```
[edit chassis]
user@PE1# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to multicast senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
```



```
user@PE1# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 59).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 59) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. At the global level and also in the bridge domain, replicate IGMP join and leave messages from the active link to the standby link of a dual-link MC-LAG interface, enabling faster recovery of membership information after failover.

```
[edit multicast-snooping-options]
user@PE1# set multichassis-lag-replicate-state

[edit bridge-domains bd0 multicast-snooping-options]
```


user@PE1# **set multichassis-lag-replicate-state**

9. Configure multicast snooping for the MC-LAG interfaces.

   ```
   [edit bridge-domains bd0]
   user@PE1# set protocols igmp-snooping vlan 100 interface ge-1/1/4.0
     multicast-router-interface
   user@PE1# set protocols igmp-snooping vlan 101 interface ge-1/1/4.0
     multicast-router-interface
   user@PE1# set protocols igmp-snooping vlan 200 interface ge-1/1/4.0
     multicast-router-interface
   ```

10. Configure ICCP parameters.

   ```
   [edit protocols iccp]
   user@PE1# set local-ip-addr 100.100.100.1
   user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
   user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
   ```

11. Configure the service ID at the global level.

   ```
   [edit switch-options]
   user@PE1# set service-id 10
   ```

   You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show multicast-snooping-options**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
  multicast-snooping-options {
    multichassis-lag-replicate-state;
  }
  protocols {
    igmp-snooping {
      vlan 100 {
        interface ge-1/1/4.0 {
          multicast-router-interface;
        }
      }
      vlan 101 {
        interface ge-1/1/4.0 {
```

```
                multicast-router-interface;
            }
        }
        vlan 200 {
            interface ge-1/1/4.0 {
                multicast-router-interface;
            }
        }
    }
  }
}
user@PE1# show chassis
aggregated-devices {
    ethernet {
        device-count 5;
    }
}
user@PE1# show interfaces
ge-1/0/1 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/0/2 {
    unit 0 {
        family inet {
            address 100.100.100.1/30;
        }
    }
}
ge-1/1/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
ge-1/1/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
```

```
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 5;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
```

```
user@PE1# show multicast-snooping-options
multichassis-lag-replicate-state;
```

```
user@PE1# show protocols
iccp {
    local-ip-addr 100.100.100.1;
    peer 100.100.100.2 {
        redundancy-group-id-list 10;
```



```
      liveness-detection {
        minimum-interval 1000;
      }
    }
  }
```

user@PE1# **show switch-options**
service-id 10;

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

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

   [edit chassis]
   user@CE# **set aggregated-devices ethernet device-count 2**

2. Specify the members to be included within the aggregated Ethernet bundle.

   [edit interfaces]
   user@CE# **set ge-2/0/2 gigether-options 802.3ad ae0**

user@CE# **set ge-2/0/3 gigether-options 802.3ad ae0**

3. Configure an interface that connects to multicast senders or receivers.

   [edit interfaces ge-2/1/6]
   user@CE# **set flexible-vlan-tagging**
   user@CE# **set encapsulation flexible-ethernet-services**
   user@CE# **set unit 0 encapsulation vlan-bridge**
   user@CE# **set unit 0 vlan-id-range 100-110**

4. Configure parameters on the aggregated Ethernet bundle.

   [edit interfaces ae0]
   user@CE# **set flexible-vlan-tagging**
   user@CE# **set encapsulation flexible-ethernet-services**
   user@CE# **set unit 0 encapsulation vlan-bridge**
   user@CE# **set unit 0 vlan-id-range 100-500**

5. Configure LACP on the aggregated Ethernet bundle.

   [edit interfaces ae0 aggregated-ether-options]
   user@CE# **set lacp active**
   user@CE# **set lacp system-priority 100**

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

   [edit bridge-domains bd0]
   user@CE# **set domain-type bridge**
   user@CE# **set vlan-id all**
   user@CE# **set interface ge-2/1/6.0**
   user@CE# **set interface ae0.0**

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}
```



```
user@CE# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}

user@CE# show interfaces
ge-2/0/2 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/3 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/1/6 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ae0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
```

```
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

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

   ```
   [edit chassis]
   user@P# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@P# set ge-1/0/5 gigether-options 802.3ad ae1
   user@P# set ge-1/0/11 gigether-options 802.3ad ae1
   ```

3. Configure an interface that connects to multicast senders or receivers.

   ```
   [edit interfaces ge-1/1/3]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-500
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-110
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1 aggregated-ether-options]
   user@P# set lacp active
   user@P# set lacp system-priority 100
   ```

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@P# set vlan-id all
   user@P# set domain-type bridge
   ```

```
user@P# set interface ge-1/1/3.0
user@P# set interface ae1.0
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-1/1/3.0;
  interface ae1.0;
}
```

```
user@P# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@P# show interfaces
ge-1/0/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/11 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/1/3 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
```

```
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- ***show igmp snooping interface***
- ***show igmp snooping membership***
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show l2-learning instance extensive**
- ***show multicast snooping route* extensive**

Related Documentation

- IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview on page 20
- Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers on page 130
- *Configuring ICCP for MC-LAG*
- **show interfaces (Aggregated Ethernet) on page 265** in the CLI Explorer

Configuring Aggregated Ethernet Link Protection

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

NOTE: Link protection is not supported on MX80.

- Configuring Link Protection for Aggregated Ethernet Interfaces on page 147
- Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 147

Configuring Link Protection for Aggregated Ethernet Interfaces

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

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

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

2. Configure the link protection mode.

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

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* primary**

2. Configure the backup logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* backup**

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

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

Disabling Link Protection for Aggregated Ethernet Interfaces

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

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

Example: Configuring Aggregated Ethernet Link Protection

The following configuration enables link protection on the **ae0** interface, and specifies the **ge-1/0/0** interface as the primary link and **ge-1/0/1** as the secondary link.

```
[edit interfaces]
ae0 {
  aggregated-ether-options {
    link protection;
  }
}
[edit interfaces]
ge-1/0/0 {
  gigether-options {
    802.3ad ae0 primary;
  }
}
[edit interfaces]
ge-1/0/1 {
  gigether-options {
    802.3ad ae0 backup;
  }
}
```

Related Documentation

- *Ethernet Interfaces*

Configuring Adaptive Load Balancing

This topic describes how to configure adaptive load balancing on PTX Series Packet Transport Routers. Adaptive load balancing maintains efficient utilization of member link bandwidth for an aggregated Ethernet (AE) bundle. Adaptive load balancing uses a feedback mechanism to correct traffic load imbalance by adjusting the bandwidth and packet streams on links within an AE bundle.

Before you begin:

- Configure a set of router interfaces with a protocol family and IP address. These interfaces can make up the membership for the AE bundle.
- Create an AE bundle by configuring a set of router interfaces as aggregated Ethernet and with a specific AE group identifier.

To configure adaptive load balancing for an AE bundle on PTX Series routers:

1. Enable adaptive load balancing on the AE bundle:

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

2. Configure the scan interval value for adaptive load balancing on the AE bundle. The scan interval value determines the length of the traffic scan by multiplying the integer value with a 30-second time period:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set scan-interval multiplier
```

3. Configure the tolerance percentage value. The tolerance value determines the allowed deviation in the traffic rates among the members of the AE bundle before the router triggers an adaptive load balancing update:

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

4. (Optional) Enable packet-per-second-based adaptive load balancing on the AE bundle:

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

Related Documentation

- Understanding Aggregated Ethernet Load Balancing on page 26
- Example: Configuring Aggregated Ethernet Load Balancing on page 151
- **adaptive on page 217**

Example: Configuring Aggregated Ethernet Load Balancing

Understanding Aggregated Ethernet Load Balancing

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

  Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

  By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

  ECMP load balancing can be:

  - Across BGP paths (BGP multipath)
  - Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

  Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

  To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  To configure the tolerance value as a percentage, include the **tolerance** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the **pps** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the **scan-interval** optional keyword at the **[edit interfaces ae*x* aggregated-ether-options load-balance adaptive]** hierarchy level.

  NOTE: The **pps** and **scan-interval** optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-Packet Random Spray Load Balancing

  When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

NOTE: **The Per-Packet option for load balancing is not supported on PTX Series Packet Transport Routers.**

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the **show interfaces ae*x* aggregated-ether-options load-balance** command.

Example: Configuring Aggregated Ethernet Load Balancing

This example shows how to configure aggregated Ethernet load balancing.

Requirements

This example uses the following hardware and software components:

- Three MX Series routers with MIC and MPC interfaces or three PTX Series Packet Transport Routers with PIC and FPC interfaces
- Junos OS Release 13.3 or later running on all devices

Overview

Load balancing is required on the forwarding plane when there are multiple paths or interfaces available to the next hop router, and it is best if the incoming traffic is load balanced across all available paths for better link utilization.

Aggregated Ethernet bundle is a typical application that uses load balancing to balance traffic flows across the member links of the bundle (IEEE 802.3ad).

Starting with Junos OS Release 13.3, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on MICs or MPCs of MX Series routers. Starting with Junos OS Release 14.1, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on PICs or FPCs of PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are:

- Adaptive—Adaptive load balancing is used in scenarios where flow-based hashing is not sufficient to achieve a uniform load distribution. This load-balancing solution implements a real-time feedback and control mechanism to monitor and manage imbalances in network load.

  The adaptive load-balancing solution corrects the traffic flow imbalance by modifying the selector entries, and periodically scanning the link utilization on each member link of the AE bundle to detect any deviations. When a deviation is detected, an adjustment event is triggered and fewer flows are mapped to the affected member link. As a result, the offered bandwidth of that member link goes down. This causes a continuous feedback loop, which over a period of time ensures that the same amount of byte rate is offered to all the member links, thus providing efficient traffic distribution across each member link in the AE bundle.

  To configure adaptive load balancing, include the **adaptive** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

  The **pps** option enables load balancing based on the packets-per-second rate. The default setting is bits-per-second load balancing.

  The **scan-interval** value configures the length of time for scanning as a multiple of 30 seconds.

  The **tolerance** value is the limit to the variance in the packet traffic flow to the aggregated Ethernet links in the bundle. You can specify a maximum of 100-percent variance. When the tolerance attribute is not configured, a default value of 20 percent is enabled for adaptive load balancing. A smaller tolerance value balances better bandwidth, but takes a longer convergence time.

  NOTE: The pps and scan-interval optional keywords are supported on PTX Series Packet Transport Routers only.

- Per-packet random spray—When the adaptive load-balancing solution fails, per-packet random spray acts as a last resort. The per-packet random spray load-balancing solution helps to address traffic imbalance by randomly spraying the packets to the aggregate next hops. This ensures that all the member links of the AE bundle are equally loaded, resulting in packet reordering.

  In addition, per-packet random spray identifies the ingress Packet Forwarding Engine that caused the traffic imbalance and eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  To configure per-packet random spray load balancing, include the **per-packet** statement at the **[edit interfaces ae*x* aggregated-ether-options load-balance]** hierarchy level.

NOTE: The Per-Packet option for load balancing is not supported on the PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being implemented by issuing the **show interfaces aex aggregated-ether-options load-balance** command.

Topology

In this topology, two aggregated Ethernet bundles - ae0 and ae1 - are configured on the links between the R2 and R3 routers.

Figure 19: Aggregated Ethernet Load Balancing

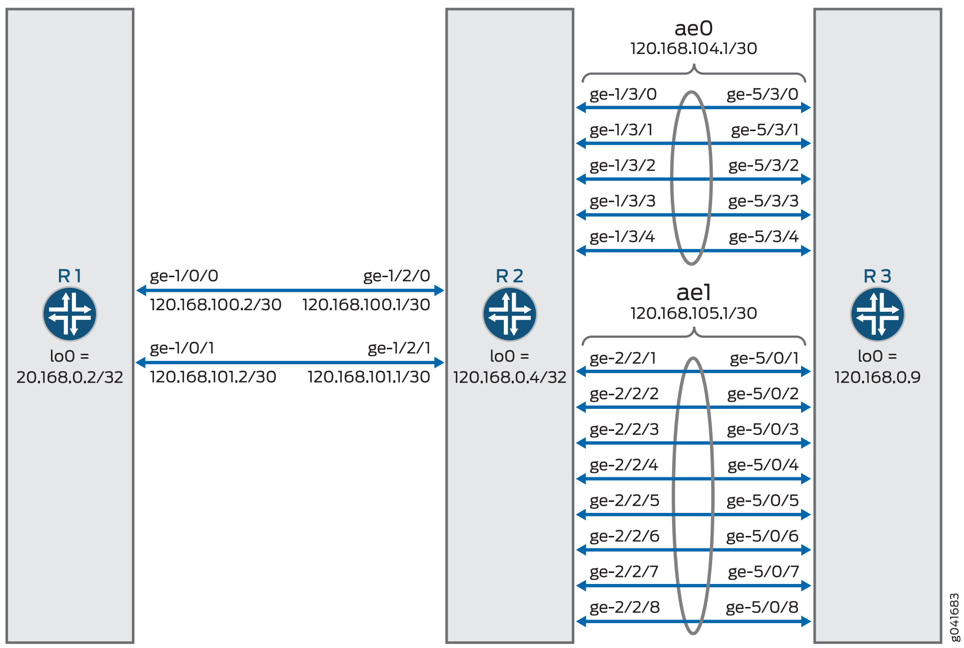

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

R1

```
set chassis aggregated-devices ethernet device-count 12
set interfaces xe-0/0/0 unit 0 family inet address 120.168.1.1/30
set interfaces xe-0/0/0 unit 0 family iso
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/1 unit 0 family inet address 120.168.2.1/30
set interfaces xe-0/0/1 unit 0 family iso
set interfaces xe-0/0/1 unit 0 family mpls
set interfaces ge-1/0/0 unit 0 family inet address 120.168.100.2/30
```



```
set interfaces ge-1/0/0 unit 0 family iso
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 family inet address 120.168.101.2/30
set interfaces ge-1/0/1 unit 0 family iso
set interfaces ge-1/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0002.00
set routing-options router-id 120.168.0.2
set routing-options autonomous-system 55
set protocols rsvp interface ge-1/0/0.0
set protocols rsvp interface ge-1/0/1.0
set protocols mpls label-switched-path videl-to-sweets to 120.168.0.9
set protocols mpls label-switched-path v-2-s-601 to 60.0.1.0
set protocols mpls label-switched-path v-2-s-601 primary v-2-s-601-primary hop-limit
  5
set protocols mpls label-switched-path v-2-s-602 to 60.0.2.0
set protocols mpls label-switched-path v-2-s-602 primary v-2-s-602-primary hop-limit
  5
set protocols mpls label-switched-path v-2-s-603 to 60.0.3.0
set protocols mpls label-switched-path v-2-s-604 to 60.0.4.0
set protocols mpls path v-2-s-601-primary 120.168.100.1 strict
set protocols mpls path v-2-s-601-primary 120.168.104.2 strict
set protocols mpls path v-2-s-602-primary 120.168.101.1 strict
set protocols mpls path v-2-s-602-primary 120.168.105.2 strict
set protocols mpls interface ge-1/0/0.0
set protocols mpls interface ge-1/0/1.0
set protocols mpls interface xe-0/0/1.0
set protocols mpls interface xe-0/0/0.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.2
set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.9
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/0/0.0
set protocols isis interface ge-1/0/1.0
set protocols isis interface lo0.0
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add
  vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-0/0/0.0
set routing-instances vpn-m5 interface xe-0/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.2:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
```



```
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.1.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.2.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-0/0/1.0
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-0/0/0.0
```

R2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/2/0 unit 0 family inet address 120.168.100.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 120.168.101.1/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/3/0 gigether-options 802.3ad ae0
set interfaces ge-1/3/1 gigether-options 802.3ad ae0
set interfaces ge-1/3/2 gigether-options 802.3ad ae0
set interfaces ge-1/3/3 gigether-options 802.3ad ae0
set interfaces ge-1/3/4 gigether-options 802.3ad ae0
set interfaces ge-2/2/1 gigether-options 802.3ad ae1
set interfaces ge-2/2/2 gigether-options 802.3ad ae1
set interfaces ge-2/2/3 gigether-options 802.3ad ae1
set interfaces ge-2/2/4 gigether-options 802.3ad ae1
set interfaces ge-2/2/5 gigether-options 802.3ad ae1
set interfaces ge-2/2/6 gigether-options 802.3ad ae1
set interfaces ge-2/2/7 gigether-options 802.3ad ae1
set interfaces ge-2/2/8 gigether-options 802.3ad ae1
set interfaces ae0 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.1/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.1/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
set accounting-options selective-aggregate-interface-stats disable
set protocols rsvp interface ge-1/2/0.0
set protocols rsvp interface ge-1/2/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/2/0.0
```

```
set protocols isis interface ge-1/2/1.0
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0
```

R3

```
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-4/0/0 unit 0 family inet address 120.168.9.1/30
set interfaces xe-4/0/0 unit 0 family mpls
set interfaces xe-4/0/1 unit 0 family inet address 120.168.10.1/30
set interfaces xe-4/0/1 unit 0 family mpls
set interfaces ge-5/0/1 gigether-options 802.3ad ae1
set interfaces ge-5/0/2 gigether-options 802.3ad ae1
set interfaces ge-5/0/3 gigether-options 802.3ad ae1
set interfaces ge-5/0/4 gigether-options 802.3ad ae1
set interfaces ge-5/0/5 gigether-options 802.3ad ae1
set interfaces ge-5/0/6 gigether-options 802.3ad ae1
set interfaces ge-5/0/7 gigether-options 802.3ad ae1
set interfaces ge-5/0/8 gigether-options 802.3ad ae1
set interfaces ge-5/3/0 gigether-options 802.3ad ae0
set interfaces ge-5/3/1 gigether-options 802.3ad ae0
set interfaces ge-5/3/2 gigether-options 802.3ad ae0
set interfaces ge-5/3/3 gigether-options 802.3ad ae0
set interfaces ge-5/3/4 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.2/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.2/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.9/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0009.00
set routing-options router-id 120.168.0.9
set routing-options autonomous-system 55
set protocols rsvp interface xe-4/0/0.0
set protocols rsvp interface xe-4/0/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls label-switched-path to-videl to 120.168.0.2
set protocols mpls interface xe-4/0/0.0
set protocols mpls interface xe-4/0/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.9
set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.2
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0
```



```
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add
  vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from protocol direct
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-4/0/0.0
set routing-instances vpn-m5 interface xe-4/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.9:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.9.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.10.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/0.0
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/1.0
```

Configuring Adaptive Load Balancing

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the R2 router:

NOTE: Repeat this procedure for the other routers, after modifying the appropriate interface names, addresses, and any other parameters for each router.

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

   ```
   [edit chassis]
   user@R2# set aggregated-devices ethernet device-count 5
   ```

2. Configure the Gigabit Ethernet interface link connecting R2 to R1.

   ```
   [edit interfaces]
   user@R2# set ge-1/2/0 unit 0 family inet address 120.168.100.1/30
   user@R2# set ge-1/2/0 unit 0 family iso
   user@R2# set ge-1/2/0 unit 0 family mpls
   ```

```
user@R2# set ge-1/2/1 unit 0 family inet address 120.168.101.1/30
user@R2# set ge-1/2/1 unit 0 family iso
user@R2# set ge-1/2/1 unit 0 family mpls

user@R2# set lo0 unit 0 family inet address 120.168.0.4/32
user@R2# set lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
```

3. Configure the five member links of the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ge-1/3/0 gigether-options 802.3ad ae0
user@R2# set ge-1/3/1 gigether-options 802.3ad ae0
user@R2# set ge-1/3/2 gigether-options 802.3ad ae0
user@R2# set ge-1/3/3 gigether-options 802.3ad ae0
user@R2# set ge-1/3/4 gigether-options 802.3ad ae0
```

4. Configure the eight member links of the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ge-2/2/1 gigether-options 802.3ad ae1
user@R2# set ge-2/2/2 gigether-options 802.3ad ae1
user@R2# set ge-2/2/3 gigether-options 802.3ad ae1
user@R2# set ge-2/2/4 gigether-options 802.3ad ae1
user@R2# set ge-2/2/5 gigether-options 802.3ad ae1
user@R2# set ge-2/2/6 gigether-options 802.3ad ae1
user@R2# set ge-2/2/7 gigether-options 802.3ad ae1
user@R2# set ge-2/2/8 gigether-options 802.3ad ae1
```

5. Enable aggregate Ethernet load balancing on ae0 of R2.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options load-balance adaptive tolerance 10
```

6. Configure the link speed for the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options link-speed 1g
```

7. Configure LACP on the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options lacp active
```

8. Configure the interface parameters for the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 unit 0 family inet address 120.168.104.1/30
user@R2# set ae0 unit 0 family iso
user@R2# set ae0 unit 0 family mpls
```

9. Enable aggregate Ethernet load balancing on ae1 of R2.

```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options load-balance adaptive tolerance 10
```

10. Configure the link speed for the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options link-speed 1g
```

11. Configure LACP on the ae1 aggregated Ethernet bundle.


```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options lacp active
```

12. Configure the interface parameters for the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae1 unit 0 family inet address 120.168.105.1/30
user@R2# set ae1 unit 0 family iso
user@R2# set ae1 unit 0 family mpls
```

13. Disable selective aggregate Ethernet statistics.

```
[edit accounting-options]
user@R2# set selective-aggregate-interface-stats disable
```

14. Configure RSVP on all the interfaces of R2 and on the AE bundles.

```
[edit protocols]
user@R2# set rsvp interface ge-1/2/0.0
user@R2# set rsvp interface ge-1/2/1.0
user@R2# set rsvp interface ae0.0
user@R2# set rsvp interface ae1.0
```

15. Configure MPLS on all the interfaces of R2 and on the AE bundles.

```
[edit protocols]
user@R2# set mpls interface ge-1/2/0.0
user@R2# set mpls interface ge-1/2/1.0
user@R2# set mpls interface ae0.0
user@R2# set mpls interface ae1.0
```

16. Configure IS-IS on all the interfaces of R2 and on the AE bundles.

```
[edit protocols]
user@R2# set isis traffic-engineering family inet shortcuts
user@R2# set isis level 1 disable
user@R2# set isis interface ge-1/2/0.0
user@R2# set isis interface ge-1/2/1.0
user@R2# set isis interface ae0.0
user@R2# set isis interface ae1.0
user@R2# set isis interface lo0.0
```

Results

From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show accounting-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
  }
}

user@R2# show interfaces
ge-1/2/0 {
  unit 0 {
```



```
        family inet {
            address 120.168.100.1/30;
        }
        family iso;
        family mpls;
    }
}
ge-1/2/1 {
    unit 0 {
        family inet {
            address 120.168.101.1/30;
        }
        family iso;
        family mpls;
    }
}
ge-1/3/0 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/1 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/2 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/3 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/4 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/2/1 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/2 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/3 {
    gigether-options {
        802.3ad ae1;
    }
}
```



```
ge-2/2/4 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/6 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/7 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/8 {
  gigether-options {
    802.3ad ae1;
  }
}
ae0 {
  aggregated-ether-options {
    load-balance {
      adaptive tolerance 10;
    }
    link-speed 1g;
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 120.168.104.1/30;
    }
    family iso;
    family mpls;
  }
}
ae1 {
  aggregated-ether-options {
    load-balance {
      adaptive tolerance 10;
    }
    link-speed 1g;
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 120.168.105.1/30;
```

```
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 120.168.0.4/32;
        }
        family iso {
            address 49.0001.1201.6800.0004.00;
        }
    }
}
```

```
user@R2# show accounting-options
selective-aggregate-interface-stats disable;
```

```
user@R2# show protocols
rsvp {
    interface ge-1/2/0.0;
    interface ge-1/2/1.0;
    interface ae0.0;
    interface ae1.0;
}
mpls {
    interface ge-1/2/0.0;
    interface ge-1/2/1.0;
    interface ae0.0;
    interface ae1.0;
}
isis {
    traffic-engineering {
        family inet {
            shortcuts;
        }
    }
    level 1 disable;
    interface ge-1/2/0.0;
    interface ge-1/2/1.0;
    interface ae0.0;
    interface ae1.0;
    interface lo0.0;
}
```

Verification

Confirm that the configuration is working properly.

- Verifying Adaptive Load Balancing on ae0 on page 162

Verifying Adaptive Load Balancing on ae0

Purpose Verify that packets received on the ae0 aggregated Ethernet bundle are load-balanced among the five member links.

Action From operational mode, run the **show interfaces ae0 extensive** command.

```
user@R2> show interfaces ae0 extensive
Logical interface ae0.0 (Index 325) (SNMP ifIndex 917) (Generation 134)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :        848761          9      81247024         7616
        Output:  166067308909    3503173 126900990064983  21423804256
    Adaptive Statistics:
        Adaptive Adjusts:        264
        Adaptive Scans  :      27682
        Adaptive Updates:         10
    Link:
      ge-1/3/0.0
        Input :        290888          5      29454436         3072
        Output:   33183442699     704569 25358563587277   4306031760
      ge-1/3/1.0
        Input :        162703          1      14806325          992
        Output:   33248375409     705446 25406995966732   4315342152
      ge-1/3/2.0
        Input :        127448          1      12130566          992
        Output:   33184552729     697572 25354827700261   4267192376
      ge-1/3/3.0
        Input :        121044          1      11481262         1280
        Output:   33245875402     697716 25405953405192   4265750584
      ge-1/3/4.0
        Input :        146678          1      13374435         1280
        Output:   33205071207     697870 25374651121458   4269487384
```

Meaning The member links of the ae0 aggregated Ethernet bundle are fully utilized with adaptive load balancing.

Configuring Shared Scheduling on Aggregated Ethernet Interfaces

You can configure shared scheduling on aggregated Ethernet interfaces in link protection mode on Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Ethernet Enhanced IQ2 (IQ2E) PICs on M320 routers.

To configure shared scheduling on aggregated Ethernet interfaces:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

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

2. Configure the link protection mode.

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

3. Configure shared scheduling.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **top**
   [edit]
   user@host# **edit** interfaces **ae*x* shared-scheduler**

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 146

Configuring the Number of Aggregated Ethernet Interfaces on the Device

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them.

On M Series and T Series routers, you can configure a maximum number of 128 aggregated interfaces, whereas on MX Series routers you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces are numbered from **ae0** through **ae127** for M Series and T Series routers and the aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers.

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

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

2. Set the number of aggregated Ethernet interfaces.

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

You must also specify the constituent physical links by including the **802.3ad** statement at the **[edit interfaces *interface-name* fastether-options]** or **[edit interfaces *interface-name* gigether-options]** hierarchy level.

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

Related Documentation

- For information about physical links, see Configuring an Aggregated Ethernet Interface on page 34
- For a sample configuration, see Example: Configuring Aggregated Ethernet Interfaces on page 35
- *Ethernet Interfaces*
- For information about configuring aggregated devices, see the *Junos OS Administration Library for Routing Devices*.

Configuring Aggregated Ethernet LACP

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

NOTE: **Starting with Junos OS Release 14.1, you can configure aggregated Ethernet interfaces with LACP on logical systems within an MX Series router.**

For Multichassis Link Aggregation (MC-LAG), you must specify the **system-id** and **admin key**. MC-LAG peers use the same **system-id** while sending the LACP messages. The **system-id** can be configured on the MC-LAG network device and synchronized between peers for validation.

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

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

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

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

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

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

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

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

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

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

Starting with Junos OS Release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.

To override the IEEE 802.3ad standard and facilitate subsecond failover, include the **fast-failover** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you use the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

For more information, see the following sections:

Configuring the LACP Interval

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options lacp]
periodic interval;
```

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

NOTE: Source address filtering does not work when LACP is enabled. This behavior is not applicable to T Series routers and PTX Series Packet Transport Routers. For more information about source address filtering, see *Enabling Ethernet MAC Address Filtering*.

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the *Routing Policy Feature Guide for Routing Devices*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see "Configuring Untagged Aggregated Ethernet Interfaces" on page 172.

For M Series Multiservice Edge Routers with enhanced Flexible PIC Concentrators (FPCs) and T Series routers, LACP over VLAN-tagged aggregated Ethernet interfaces is supported. For 8-port, 12-port, and 48-port Fast Ethernet PICs, LACP over VLAN-tagged interfaces is not supported.

LACP **Fast Periodic**, which is achieved by configuring fast (every second) intervals for periodic transmission of LACP packets, is supported with graceful Routing Engine switchover (GRES) on MX Series routers only.

Configuring LACP Link Protection

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

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the **link-protection** and **system-priority** statements. Configuring values at this level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

NOTE: LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

Enabling LACP Link Protection

To enable LACP link protection for an aggregated Ethernet interface, use the **link-protection** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable;
  revertive;
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, once a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch and the current link remains active.

If LACP link protection is configured to be nonrevertive at the global (**[edit chassis]** hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.

CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

We strongly recommend that you use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces of MX Series routers to any other vendor device. Otherwise, the vendor device (say a Layer 2 switch, or a router) will not be able to manage the traffic coming from the two link aggregated Ethernet bundle. As a result, you might observe the vendor device sending back the traffic to the backup member link of the aggregated Ethernet interface.

Currently, MX-MPC2-3D, MX-MPC2-3D-Q, MX-MPC2-3D-EQ, MX-MPC1-3D, MX-MPC1-3D-Q, and MPC-3D-16XGE-SFPP do not drop traffic coming back to the backup link, whereas DPCE-R-Q-20GE-2XGE, DPCE-R-Q-20GE-SFP, DPCE-R-Q-40GE-SFP, DPCE-R-Q-4XGE-XFP, DPCE-X-Q-40GE-SFP, and DPCE-X-Q-4XGE-XFP drop traffic coming to the backup link.

Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 to 65,535.

Configuring LACP System Identifier

To configure the LACP system identifier for aggregated Ethernet interfaces, use the **system-id** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-id system-id;
```

You must not configure the LACP system identifier by using the **system-id *system-id*** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level to be all zeros (00:00:00:00:00:00). If you attempt to commit a configuration with the system identifier to be all zeros, an error occurs during the commit operation.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Configuring LACP administrative Key

To configure an administrative key for LACP, include the **admin-key *number*** statement at the **[edit interfaces ae*x* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces ae x aggregated-ether-options-lacp]
admin-key number;
```

NOTE: **You must configure MC-LAG to configure the admin-key statement. For more information about MC-LAG, see** "Configuring Multichassis Link Aggregation" on page 42.

Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the **port-priority** statement at the **[edit interfaces *interface-name* gigether-options 802.3ad ae*X* lacp]** or **[edit interfaces *interface-name* fastether-options 802.3ad ae*X* lacp]** hierarchy levels:

```
[edit interfaces interface-name gigether-options 802.3ad aeX lacp]
port-priority priority;
```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and is assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.

NOTE: Port aggregation selection (discussed previously) is performed for the active link when LACP link protection is enabled. Without LACP link protection, port priority is not used in port aggregation selection.

Tracing LACP Operations

To trace the operations of the LACP process, include the **traceoptions** statement at the **[edit protocols lacp]** hierarchy level:

```
[edit protocols lacp]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

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

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

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

LACP Limitations

LACP can link together multiple different physical interfaces, but only features that are supported across all of the linked devices will be supported in the resulting link aggregation group (LAG) bundle. For example, different PICs can support a different number of forwarding classes. If you use link aggregation to link together the ports of a PIC that supports up to 16 forwarding classes with a PIC that supports up to 8 forwarding classes, the resulting LAG bundle will only support up to 8 forwarding classes. Similarly, linking together a PIC that supports WRED with a PIC that does not support it will result in a LAG bundle that does not support WRED.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

LACP with VLAN-Tagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}
```

Configure aggregated Ethernet LACP over an untagged interface:

LACP with Untagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}
```

Related Documentation

- lacp on page 231
- link-protection on page 233

- traceoptions on page 253
- *Ethernet Interfaces*

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces ae*x*]** hierarchy level:

```
[edit interfaces aex]
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

You can include this statement at the following hierarchy levels:

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

For more information about the **vlan-tagging** and **vlan-id** statements, see *802.1Q VLANs Overview*.

Related Documentation

- vlan-id on page 254
- vlan-tagging on page 255

Configuring Untagged Aggregated Ethernet Interfaces

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

- You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.
- You cannot include the **vlan-id** statement in the configuration of the logical interface.

Table 4 on page 172 lists untagged aggregated Ethernet and LACP support by PIC and router.

Table 4: Untagged Aggregated Ethernet and LACP Support by PIC and Platform

PIC Type	M Series	LACP	T Series	LACP
4-port Fast Ethernet PIC Type 1	Yes	Yes	Yes	Yes
1-port Gigabit Ethernet PIC Type 1	Yes	Yes	Yes	Yes

Table 4: Untagged Aggregated Ethernet and LACP Support by PIC and Platform *(continued)*

PIC Type	M Series	LACP	T Series	LACP
2-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
4-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
1-port 10-Gigabit Ethernet M160	Yes	Yes	NA	NA
10-port Gigabit Ethernet PIC Type 3	Yes (M120, M320)	Yes	Yes	Yes
1-port 10-Gigabit Ethernet PIC Type 3	N/A	NA	Yes	Yes
8-port Gigabit Ethernet PIC Type 3	Yes	Yes	Yes	Yes

The 8–port Fast Ethernet PIC does not support untagged aggregated Ethernet or LACP.

Syslog messages are logged if you try to configure an untagged aggregated Ethernet interface using an unsupported PIC type.

For more information about configuring LACP, see “Configuring Aggregated Ethernet LACP” on page 164.

Example: Configuring Untagged Aggregated Ethernet Interfaces

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

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

Related Documentation

- For more information about configuring LACP, see Configuring Aggregated Ethernet LACP on page 164.

- *Ethernet Interfaces*

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

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

NOTE:

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

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

To set the required link speed:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces *interface-name*** aggregated-ether-options

2. Configure the link speed.

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

   speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet interfaces on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.
- **50g**—Links are 50 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

Related Documentation

- **aggregated-ether-options on page 219**
- Configuring Mixed Aggregated Ethernet Links on page 41

- *Ethernet Interfaces*

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces** ***interface-name*** aggregated-ether-options

2. Configure the minimum number of links.

   [edit interfaces *interface-name* aggregated-ether-options]
   user@host# **set** minimum-links ***number***

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links** ***number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links** ***number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links** ***number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

Related Documentation

- aggregated-ether-options on page 219
- minimum-links on page 240

- *Ethernet Interfaces*

Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces

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

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

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

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

To configure multicast statistics:

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

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

```
[edit interfaces]
ae0 {
  multicast-statistics;
}
```

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

Related Documentation

- *multicast-statistics*
- *Configuring Multicast Statistics Collection on Ethernet Interfaces*
- *Ethernet Interfaces*

Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection

On aggregated Ethernet interfaces, you can configure scheduler in non–link-protect mode on the following platforms:

- MX-Series
- M120 and M320 with IQ2 PIC
- T-series platforms (T620 and T320) with IQ2 PIC

The scheduler functions supported are:

- Per unit scheduler
- Hierarchical scheduler
- Shaping at the physical interface

To configure the hierarchical scheduler on aggregated Ethernet interfaces in the non link-protect mode, include the **hierarchical-scheduler** statement at the **[edit interfaces ae*X*]** hierarchy level:

```
[edit interfaces aeX hierarchical-scheduler]
```

Prior to Junos OS Release 9.6, the hierarchical scheduler mode on these models required the **aggregated-ether-options** statement **link-protection** option. If a **link-protection** option is not specified, the scheduler is configured in non-link-protect mode.

To specify the member link bandwidth derivation based on the equal division model **(scale)** or the replication model **(replicate)** on aggregated Ethernet interfaces, include the **member-link-scheduler (scale | replicate)** option at the **[edit class-of-service interfaces ae*X*]** hierarchy level. The default setting is **scale**.

[edit class-of-service interfaces ae*X* member-link-scheduler (scale | replicate)]

NOTE: In link-protect mode, only one link is active at a time and the other link acts as the backup link, whereas in a non link-protect mode, all the links of the aggregate bundle are active at the same time. There is no backup link. If a link goes down or a new link is added to the bundle, traffic redistribution occurs.

Related Documentation

- *Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links*
- *Ethernet Interfaces*
- For more information on the hierarchical scheduler (CoS), see the *Junos OS Class of Service Library for Routing Devices*.

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

This section describes configuration of symmetrical load balancing on an 802.3ad link aggregation group (LAG) on MX Series routers.

Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview

MX Series routers with Aggregated Ethernet PICs support symmetrical load balancing on an 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through deep packet inspection (DPI) devices over an LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Without symmetrical load balancing on an 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. By using this feature, a given flow of traffic (duplex) is ensured for the same devices in both directions.

Symmetrical load balancing on an 802.3ad LAG utilizes a mechanism of interchanging the source and destination addresses for a hash computation of fields, such as source address and destination address. The result of a hash computed on these fields is used to choose the link of the LAG. The hash-computation for the forward and reverse flow must be identical. This is achieved by swapping source fields with destination fields for the reverse flow. The swapped operation is referred to as *complement hash computation* or **symmetric-hash complement** and the regular (or unswapped) operation as *symmetric-hash computation* or **symmetric-hash**. The swappable fields are MAC address, IP address, and port.

Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

You can specify whether symmetric hash or complement hash is done for load-balancing traffic. To configure symmetric hash, use the **symmetric-hash** statement at the **[edit forwarding-options hash-key family inet]** hierarchy level. To configure symmetric hash complement, use the **symmetric-hash complement** statement and option at the **[edit forwarding-options hash-key family inet]** hierarchy level.

These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the **symmetric-hash** or **symmetric-hash complement** statement at the **[edit chassis hash-key family inet]** and **[edit chassis hash-key family multiservice]** hierarchy levels.

Consider the example in Figure 20 on page 179.

Figure 20: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

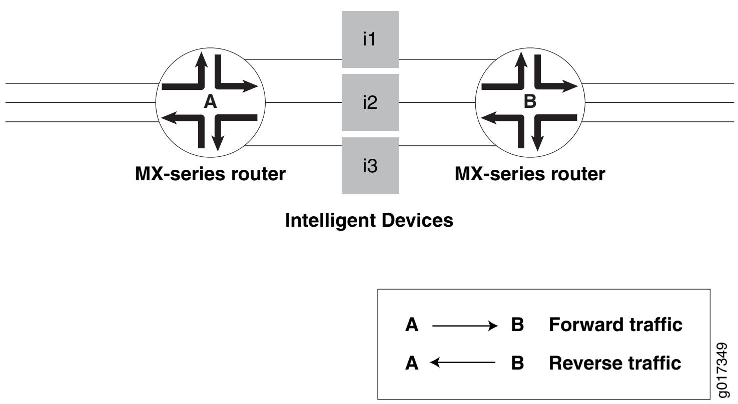

Router A is configured with symmetric hash and Router B is configured with symmetric hash complement. Thus, for a given flow *fx*, post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow *fx* is from Router B to Router A through the same i2 device as its hashing (done after swapping source and destination fields) and returns the same link index; since it is performed on the interchanged source and destination addresses.

However, the link chosen may or may not correspond to what was attached to the DPI. In other words, the hashing result should point to the same links that are connected, so that the traffic flows through the same DPI devices in both directions. To make sure this happens, you need to also configure the counterpart ports (ports that are connected to same DPI-iN) with the identical link index. This is done when configuring a child-link into the LAG bundle. This ensures that the link chosen for a given hash result is always the same on either router.

Note that any two links connected to each other should have the same link index and these link indices must be unique in a given bundle.

NOTE:

The following restrictions apply when configuring symmetric load balancing on an 802.3ad LAG on MX Series routers:

- The Packet Forwarding Engine (PFE) can be configured to hash the traffic in either symmetric or complement mode. A single PFE complex cannot work simultaneously in both operational modes and such a configuration can yield undesirable results.
- The per-PFE setting overrides the chassis-wide setting only for the family configured. For the other families, the PFE complex still inherits the chassis-wide setting (when configured) or the default setting.
- This feature supports VPLS, INET, and bridged traffic only.
- This feature cannot work in tandem with the **per-flow-hash-seed load-balancing** option. It requires that all the PFE complexes configured in complementary fashion share the same seed. A change in the seed between two counterpart PFE complexes may yield undesired results.

For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library for Routing Devices*.

Example Configuration Statements

To configure 802.3ad LAG parameters at the bundle level:

```
[edit interfaces]
g(x)e-fpc/pic/port {
   gigether-options {
      802.3ad {
         bundle;
         link-index number;
      }
   }
}
```

where the **link-index *number*** ranges from 0 through 15.

You can check the link index configured above using the **show interfaces** command:

```
[edit forwarding-options hash-key]
family inet {
   layer-3;
   layer-4;
   symmetric-hash {
      [complement;]
   }
}
family multiservice {
   source-mac;
   destination-mac;
   payload {
   ip {
      layer-3 {
         source-ip-only | destination-ip-only;
```

```
        }
        layer-4;
      }
    }
    symmetric-hash {
      [complement;]
    }
  }
```

For load-balancing Layer 2 traffic based on Layer 3 fields, you can configure 802.3ad LAG parameters at a per PIC level. These configuration options are available under the chassis hierarchy as follows:

```
[edit chassis]
fpc X {
  pic Y {
    .
    .
    .
    hash-key {
      family inet {
        layer-3;
        layer-4;
        symmetric-hash {
          [complement;]
        }
      }
      family multiservice {
        source-mac;
        destination-mac;
        payload {
          ip {
            layer-3 {
              source-ip-only | destination-ip-only;
            }
            layer-4;
          }
        }
        symmetric-hash {
          [complement;]
        }
      }
    }
    .
    .
    .
  }
}
```

Configuring Symmetrical Load Balancing on Trio-Based MPCs

With some configuration differences, symmetrical load-balancing over an 802.3ad link aggregation group is supported on MX Series routers with Trio-based MPCs.

To achieve symmetrical load-balancing on Trio-Based MPCs, the following needs to be done:

- Compute a Symmetrical Hash

  Both routers must compute the same hash value from the flow in the forward and reverse directions.On Trio-based platforms, the calculated hash value is independent of the direction of the flow, and hence is always symmetric in nature. For this reason, no specific configuration is needed to compute a symmetric hash value on Trio-based platforms.

  However, it should be noted that the fields used to configure the hash should have identical include and exclude settings on both ends of the LAG.

- Configure Link Indexes

  To allow both routers to choose the same link using the same hash value, the links within the LAG must be configured with the same link index on both routers. This can be achieved with the **link-index** statement.

- Enable Symmetric Load Balancing

  To configure symmetric load balancing on Trio-based MPCs, include the **symmetric** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy level. This statement is applicable to Trio-based platforms only.

  The **symmetric** statement can be used with any protocol family and enables symmetric load-balancing for all aggregated Ethernet bundles on the router. The statement needs to be enabled at both ends of the LAG. This statement is disabled by default.

- Achieve Symmetry for Bridged and Routed Traffic

  In some deployments, the LAG bundle on which symmetry is desired is traversed by Layer 2 bridged traffic in the upstream direction and by IPv4 routed traffic in the downstream direction. In such cases, the computed hash is different in each direction because the Ethernet MAC addresses are taken into account for bridged packets. To overcome this, you can exclude source and destination MAC addresses from the enhanced-hash-key computation.

  To exclude source and destination MAC addresses from the enhanced-hash-key computation, include the **no-mac-addresses** statement at the **[edit forwarding-options enhanced-hash-key family multiservice]** hierarchy level. This statement is disabled by default.

When symmetrical load balancing is enabled on Trio-based MPCs, keep in mind the following caveats:

- Traffic polarization is a phenomenon that occurs when using topologies that distribute traffic by using hashing of the same type. When routers are cascaded, traffic polarization can occur, and this can lead to unequal traffic distribution.

  Traffic polarization occurs when LAGs are configured on cascaded routers. For example, in Figure 21 on page 183, if a certain flow uses Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, the flow also uses Link 1 of the aggregated Ethernet bundle between Device R2 and Device R3.

Figure 21: Traffic Polarization on Cascaded Routers When Symmetrical Load Balancing in Enabled on Trio-based MPCs

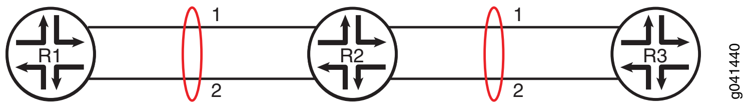

This is unlike having a random link selection algorithm, where a flow might use Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, and Link 2 of the aggregated Ethernet bundle between Device R2 and Device R3.

- Symmetric load balancing is not applicable to per-prefix load-balancing where the hash is computed based on the route prefix.
- Symmetric load balancing is not applicable to MPLS or VPLS traffic, because in these scenarios the labels are not the same in both directions.

Example Configurations

Example Configurations of Chassis Wide Settings

Router A

```
user@host> show configuration forwarding-options hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric hash;
}
```

Router B

```
user@host> show configuration forwarding-options hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric-hash {
    complement;
  }
}
```

Example Configurations of Per–Packet-Forwarding-Engine Settings

Router A

```
user@host> show configuration chassis fpc 2 pic 2 hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric hash;
}
```


Router B

```
user@host> show configuration chassis fpc 2 pic 3 hash-key
family multiservice {
   payload {
      ip {
         layer-3;
      }
   }
   symmetric-hash {
      complement;
   }
}
```

Related Documentation

- *Ethernet Interfaces*
- For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library for Routing Devices*.

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

[edit chassis] Hierarchy Level

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
    sonet {
      device-count number;
    }
  }
    maximum-links {
    }
  channel-group number {
    ethernet {
      device-count number;
    }
    fpc slot-number{
      pic pic-number {
```

```
                adaptive-services{
                  service-package (layer-2 | layer-3);
                }
                aggregate-ports;
                atm-cell-relay-accumulation;
                atm-l2circuit-mode (aal5 | cell | trunk trunk);
                ce1 {
                  e1 link-number {
                    channel-group group-number;
                    timeslots time-slot-range;
                  }
                }
                channelization;
                ct1 {
                  t1 link-number {
                    channel-group group-number;
                    timeslots time-slot-range;
                  }
                }
                ct3 {
                  port port-number {
                    t1 link-number {
                      channel-group group-number;
                      timeslots time-slot-range;
                    }
                  }
                  framing sdh;
                }
                max-queues-per-interface number;
                mlfr-uni-nni-bundles num-intf;
                no-concatenate;
                shdsl {
                  pic-mode (1-port-atm | 2-port-atm);
                }
                vtmapping (klm | itu-t);
              }
            }
            fpc slot-number{
            pic pic-number{
              account-layer2-overhead
                egress-policer-overhead bytes;
                ingress-policer-overhead bytes;
                mlfr-uni-nni-bundles-inline number;
                multi-link-layer-2-inline;
              }
            }
          }
```

Related Documentation

- *Router Interfaces*

[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 connections] Hierarchy Level

The following statements can also be configured at the **[edit logical-systems *logical-system-name* protocols connections]** hierarchy level.

```
interface-switch connection-name {
  interface interface-name.unit-number;
  interface interface-name.unit-number;
}
```

Related Documentation

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

[edit protocols dot1x] Hierarchy Level

```
dot1x {
  authenticator
    authentication-profile-name access-profile-name;
    interface interface-ids {
      maximum-requests integer;
      retries integer;
      quiet-period seconds;
      transmit-period seconds;
      reauthentication (disable | interval seconds);
      server-timeout seconds;
      supplicant (single);
      supplicant-timeout seconds;
    }
  }
}
```

Related Documentation

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

[edit protocols iccp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
iccp {
  traceoptions;
  local-ip-address ip address;
  session-establishment-hold-time value;
  authentication-key string;
  peer ip-address {
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    redundancy-group-id-list redundancy-group-id-list;
    liveness-detection;
  }
}
```

Related Documentation

- iccp on page 224
- *Configuring ICCP for MC-LAG*

[edit protocols lacp] Hierarchy Level

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>);
  }
fast-hello-issu
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *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*

[edit protocols ppp] Hierarchy Level

```
monitor-session (interface-name | all);
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable> ;
  flag flag <disable>;
}
```

Related Documentation

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

[edit protocols pppoe] Hierarchy Level

```
pppoe {
```



```
    no-send-pads-error;
    no-send-pads-ac-info
    pado-advertise;
    service-name-tables table-name {
      service service-name {
        drop;
        delay seconds;
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        max-sessions number;
        agent-specifier {
          aci circuit-id-string ari remote-id-string {
            drop;
            delay seconds;
            terminate;
            dynamic-profile profile-name;
            routing-instance routing-instance-name;
            static-interface interface-name;
          }
        }
      }
    }
    traceoptions {
      file <filename> <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
      filter {
        aci regular-expression;
        ari regular-expresion;
        service-name regular-expresion;
        underlying-interface interface-name;
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
```

[edit protocols protection-group] Hierarchy Level

```
  ethernet-ringring-name {
    east-interface {
      control-channel channel-name {
        vlan number;
      }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    west-interface {
      control-channel channel-name {
        vlan number;
      }
    }
```



```
}
```

Related Documentation

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

[edit protocols vrrp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
   vrrp {
      asymmetric-hold-time;
      delegate-processing;
      failover-delay milliseconds;
      global-advertisements-threshold advertisement-value;
      skew-timer-disable;
      startup-silent-period seconds;
      traceoptions {
         file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size maximum-file-size> <world-readable | no-world-readable>;
         flag flag;
         no-remote-trace;
         }
      version-3;
   }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit system processes] Hierarchy Level

```
dialer-services {
   disable;
}
isdn-signaling {
   disable;
   reject-incoming;
}
```

Related Documentation

- *ISDN Configuration Overview*
- *Disabling ISDN Processes*

CHAPTER 4

Statement Summary

802.3ad

Syntax

```
802.3ad {
  ae interface-number (primary | backup);
  lacp {
    port-priority;
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
primary and **backup** options added in Junos OS Release 8.3.

Description

Specify aggregated Ethernet logical interface number.

Options

ae ***interface-number***—Aggregated Ethernet logical interface number.

Range: 0 through 15

primary—For link protection configurations, specify the primary link for egress traffic.

backup—For link protection configurations, specify the backup link for egress traffic.

Required Privilege Level

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

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 34
- Configuring Aggregated Ethernet Link Protection on page 146

adaptive

Syntax

```
adaptive {
  pps;
  scan-interval multiple;
  tolerance tolerance-percentage;
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options load-balance]

Release Information Statement introduced in Junos OS Release 13.3.
Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.

Description Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.

Options **pps**—(PTX Series only) The type of traffic rate among the members of the AE bundle is measured packets per second. The default rate type is bytes per second.

scan-interval *multiple*—(PTX Series only) Scan interval, as a multiple of a 30-second interval.

Range: 1 through 5

Default: 1

tolerance *tolerance-percentage*—(MX Series and PTX Series) Limit to the variance in the packet traffic flow to the aggregated Ethernet links in a percentage.

Range: 1 through 100 percent

Default: 20 percent

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

Related Documentation

- Understanding Aggregated Ethernet Load Balancing on page 26
- Example: Configuring Aggregated Ethernet Load Balancing on page 151

aggregated-devices

Syntax

```
aggregated-devices {
  ethernet {
    device-count number;
    lacp {
      link-protection {
        non-revertive;
      }
      system-priority;
    }
  }
  sonet {
    device-count number;
    }
  maximum-links maximum-links-limit;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 7.4.
Support for LACP link protection and system priority introduced in Junos OS Release 9.3.

Description Configure properties for aggregated devices on the router.

Options The remaining statements are explained separately.

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

Related Documentation

- Configuring Junos OS for Supporting Aggregated Devices on page 37

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  ethernet-switch-profile {
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
    (mac-learn-enable | no-mac-learn-enable);
  }
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    link-protection {
      disable;
    (revertive | non-revertive);
    periodic interval;
    system-priority priority;
    system-id system-id;
  }
  link-protection;
  load-balance;
  link-speed speed;
  logical-interface-chassis-redundancy;
  logical-interface-fpc-redundancy;
  (loopback | no-loopback);
  minimum-links number;
  rebalance-periodic time hour:minute <interval hours>;
  source-address-filter {
    mac-address;
    (source-filtering | no-source-filtering);
  }
}
```


Hierarchy Level [edit interfaces ae*x*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure aggregated Ethernet-specific interface properties.

The statements are explained separately.

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

Related Documentation
- *Ethernet Interfaces Overview*

chassis

Syntax chassis { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure router chassis properties.

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

Related Documentation
- *Router Chassis Configuration Statements*

device-count

Syntax device-count *number*;

Hierarchy Level [edit chassis aggregated-devices ethernet]
[edit chassis aggregated-devices sonet]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the number of aggregated logical devices available to the router.

Options **number**—Number of aggregated logical devices available to the router.

NOTE: Starting with Junos OS Release 13.2, a maximum of 64 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces. In releases before Junos OS Release 13.2, a maximum of 16 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces.

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

Related Documentation
- Configuring Junos OS for Supporting Aggregated Devices on page 37
- *Configuring Aggregated SONET/SDH Interfaces*

disable (Link Protection)

Syntax disable;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Disable LACP link protection on the interface.

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

Related Documentation
- *Configuring Aggregated Ethernet LACP*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

ethernet (Chassis)

Syntax

```
ethernet {
  device-count number;
  lacp {
    link-protection {
      non-revertive;
    }
    system-priority;
  }
}
```

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Configure properties for Ethernet aggregated devices on the router.

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

Related Documentation

- Configuring Junos OS for Supporting Aggregated Devices on page 37
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

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*

iccp

Syntax

```
iccp {
  traceoptions; {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
      <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
  local-ip-address ip address;
  session-establishment-hold-time value;
  authentication-key string;
  peer ip-address {
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    redundancy-group-id-list redundancy-group-id-list;
    liveness-detection;
  }
}
```

Hierarchy Level

[edit protocols iccp]
[edit logical-systems *logical-system-name* protocols iccp]

Release Information

Statement introduced in Junos OS Release 10.0.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

Default

If you do not include this statement, no ICCP protocol tracing operations are performed.

Options

traceoptions—Set Interchassis Control Protocol (ICCP) tracing options.

local-ip-address—Specify the source address where the ICCP packet is routed.

session-establishment-hold-time—Specify if the chassis takes over as the master at the ICCP session.

authentication-key —Specify TCP Message Digest 5 (MD5) option for an ICCP TCP session.

peer ip-address—Specify the IP address of the peer that hosts an MC-LAG. You must configure ICCP for both peers that host the MC-LAG.

redundancy-group-id-list—Specify the redundancy groups between two ICCP peers.

liveness-detection—Specify Bidirectional Forwarding Detection (BFD) protocol options.

Required Privilege Level

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

Related Documentation

- *Configuring ICCP for MC-LAG*

igmp-snooping

Syntax

```
igmp-snooping {
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    immediate-leave;
    multicast-router-interface;
    static {
      group ip-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
  vlan vlan-id {
    immediate-leave;
    interface interface-name {
      group-limit limit;
      host-only-interface;
      immediate-leave;
      multicast-router-interface;
      static {
        group ip-address {
          source ip-address;
        }
      }
    }
    proxy {
      source-address ip-address;
    }
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
  }
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* protocols],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols]
[edit routing-instances *routing-instance-name* protocols]
[edit protocols]

Release Information

Statement introduced in Junos OS Release 8.5.

Description

Enable IGMP snooping on the router.

Default IGMP snooping is disabled on the router.

Options The statements are explained separately.

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

Related Documentation

- *Understanding IGMP Snooping*
- IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview on page 20

inner-tag-protocol-id

Syntax inner-tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.

Default If the **inner-tag-protocol-id** statement is not configured, the TPID value is 0x8100.

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

Related Documentation

- *Configuring Inner and Outer TPIDs and VLAN IDs*

inner-vlan-id

Syntax inner-vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers or 100-Gigabit Ethernet Type 5 PIC with CFP, or on Ethernet interfaces on EX Series switches, specify the VLAN ID to rewrite for the inner tag of the final packet.

You cannot include the **inner-vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement and the **inner-vlan-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **inner-vlan-id** statement you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Options *number*—VLAN ID number.

Range: 0 through 4094

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

Related Documentation
- *Configuring Inner and Outer TPIDs and VLAN IDs*

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router or switch.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

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

Related Documentation

- *Physical Interface Configuration Statements Overview*
- Configuring Aggregated Ethernet Link Protection on page 146

lacp (802.3ad)

Syntax

```
lacp {
  traceoptions {
    file lacpd;
    flag all;
  }
  ppm (centralized | distributed);
}
```

Hierarchy Level

[edit interfaces *interface-name* fastether-options 802.3ad],
[edit interfaces *interface-name* gigether-options 802.3ad]

Release Information

Statement introduced in Junos OS Release 9.3.
The **ppm (centralized | distributed)** option introduced in Junos OS Release 9.4.

Description

For aggregated Ethernet interfaces only, configure the Link Aggregation Control Protocol (LACP).

On MX and T Series routers you can specify distributed or centralized periodic packet management (PPM).

Default

If you do not specify **lacp** as either **active** or **passive**, LACP remains passive.

If you do not specify **ppm** as either **centralized** or **distributed**, PPM is distributed.

Options

- **active**—Initiate transmission of LACP packets.
- **passive**—Respond to LACP packets.
- **ppm**—Set PPM to centralized or distributed.

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 Aggregated Ethernet LACP on page 164

lacp (Aggregated Ethernet)

Syntax

```
lacp {
  (active | passive);
  admin-key key;
  accept-data;
  fast-failover;
  link-protection {
    disable;
    (revertive |non-revertive);
  }
  periodic interval;
  system-id mac-address;
  system-priority priority;
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit logical-systems *logical-system-name* interfaces ae*X* aggregated-ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
fast-failover option introduced in Junos OS Release 12.2.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Configure the Link Aggregation Control Protocol (LACP) for aggregated Ethernet interfaces only.

When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

Default

If you do not specify LACP as either **active** or **passive**, LACP remains passive.

Options

active—Initiate transmission of LACP packets.

admin-key ***number***—Specify an administrative key for the router or switch.

NOTE: You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive—Respond to LACP packets.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Aggregated Ethernet LACP on page 164
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

link-protection

Syntax

```
link–protection {
  disable;
  (revertive |non-revertive);
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit interfaces ae*x* aggregated-ether-options *lacp*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the **[edit interfaces ge-*fpc/pic/port* gigether-options 802.3ad ae*x*]** hierarchy level or the **[edit interfaces fe-*fpc/pic/port* fastether-options 802.3ad ae*x*]** hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify **link-protection** at the following hierarchy levels:

- **[edit interfaces ge-*fpc/pic/port* ether-options 802.3ad ae*x*]**
- **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]**

For LACP standards link protection, specify **link-protection** at the following hierarchy levels:

- For global LACP link protection, specify at **[edit chassis aggregated-devices ethernet lacp]**
- For a specific aggregated Ethernet interface, specify at **[edit interfaces aeX aggregated-ether-options lacp]**

To disable link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the **[edit interfaces ae*x* aggregated-ether-options]** hierarchy level or the **[edit interfaces ae*x* aggregated-ether-options *lacp*]]** hierarchy level.

Options

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 146
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-sonet-options]

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

Description For aggregated Ethernet interfaces only, set the required link speed.

Options *speed*—For aggregated Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.

- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

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

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 3
- Configuring Aggregated Ethernet Link Speed on page 174
- Configuring Mixed Aggregated Ethernet Links on page 41
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

link-speed (Aggregated SONET/SDH)

Syntax link-speed (*speed* | mixed);

Hierarchy Level [edit interfaces as*x* aggregated-sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.
mixed option added in Release 8.0.

Description For aggregated SONET/SDH interfaces only, set the required link speed.

Options ***speed***—Aggregated SONET/SDH links can have one of the following speed values.

- **oc3**—Links are OC3c or STM1c.
- **oc12**—Links are OC12c or STM4c.
- **oc48**—Links are OC48c or STM16c.
- **oc192**—Links are OC192c or STM64c.
- **oc768**—Links are OC768c or STM256c.

mixed—For aggregated SONET/SDH links on T Series routers, you can mix interface speeds in SONET/SDH aggregation bundles. Interface speeds from OC3 through OC768 are supported.

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

Related Documentation

- Configuring Aggregated Ethernet Link Speed on page 174
- *Configuring Aggregated SONET/SDH Interfaces*

load-balance

Syntax

```
load-balance {
  adaptive{
    pps;
    scan-interval multiple;
    tolerance percentage;
  }
  no-adaptive;
  per-packet;
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 13.3.
Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.

Description Load-balances the received traffic across all the available paths of aggregated Ethernet bundles for better link utilization.

Options

adaptive— (MX Series and PTX Series) Corrects a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an Aggregated Ethernet bundle.

no-adaptive— (MX Series and PTX Series) Disables the adaptive load-balancing solution configured to distribute traffic by using a feedback mechanism.

per-packet— (MX Series only) Randomly sprays packets to the aggregate next hops in a round-robin manner to avoid traffic imbalance.

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

Related Documentation

- Understanding Aggregated Ethernet Load Balancing on page 26
- Example: Configuring Aggregated Ethernet Load Balancing on page 151

mc-ae

Syntax

```
mc-ae {
  chassis-id chassis-id;
  events {
    iccp-peer-down {
      force-icl-down;
      prefer-status-control-active;
    }
  }
  mc-ae-id mc-ae-id;
  mode (active-active | active-standby);
  redundancy-group group-id;
  revert-time revert-time;
  status-control (active | standby);
  switchover-mode (non-revertive |revertive);
}
```

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options]
[edit logical-systems *logical-system-name interfaces* ae*X* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 9.6.
events statement introduced in Junos OS Release 11.4R4.
switchover-mode and **revert-time** statements introduced in Junos OS Release 13.3.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Enable multichassis link aggregation (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

Options

chassis-id—Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links.

Values: 0 | 1

events—Specify an action if a specific MC-LAG event occurs.

iccp-peer-down—Specify an action if the ICCP peer of this node goes down.

force-icl-down—If the node's ICCP peer goes down, bring down the interchassis-link logical interface.

prefer-status-control-active—If the node's peer goes down, prefer that the node configured as **status-control active** become the active node.

NOTE: The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP does not go down unless the router is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP

BFD timeout. This configuration prevents traffic loss by ensuring that when the router with the **status-control active** configuration goes down, the router with the **status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lacp periodic** ***interval*** statement at the **[edit interface** ***interface-name*** **aggregated-ether-options]** hierarchy level as **slow** or configure the **detection-time threshold** statement at the **[edit protocols iccp liveness-detection]** hierarchy level as less than 3 seconds.

mc-ae-id—Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

Range: 1 through 65,535

mode (active-active | active-standby)—Specify whether the MC-LAG is in active-active or active-standby mode.

NOTE: You can configure IPv4 (inet) and IPv6 (inet6) addresses on **mc-ae** interfaces when the **active-standy** mode is configured.

redundancy-group—Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions.

Range: 1 through 4,294,967,294

revert-time—Wait interval (in minutes) before the switchover to the preferred node is performed when the **switchover-mode** is configured as revertive.

Range: 1 through 10

status-control (active | standby)—Specify whether the chassis becomes active or remains in standby when an interchassis link failure occurs.

switchover-mode (non-revertive | revertive)—Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.

NOTE: For **revertive** mode to automatically switch over to the preferred node, the **status-control** statement should be configured as **active**.

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

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 8
- Configuring Multichassis Link Aggregation on page 42
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 56
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on page 127

maximum-links

Syntax maximum-links *maximum-links-limit*;

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced in Junos OS Release 11.1 for T Series routers.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for MX Series routers.

Description Configure the maximum links limit for aggregated devices. Note that for MX Series routers, the router must be running in Enhanced IP mode, which restricts the chassis to running Trio-based MPCs and multiservice DPCs (MS-DPCs). For MX series routers and PTX series switches, the option for 64 links is only supported for Junos OS release 12.3 and later.

See *network-services* for more information on running in Enhanced IP mode, and *Network Services Mode Overview* for MPC specific-details and an overview of Network Services Mode.

Options *maximum-links-limit*—Maximum links limit for aggregated devices.

Range: 16, 32, 64

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

Related Documentation

- Configuring Junos OS for Supporting Aggregated Devices on page 37
- Configuring an Aggregated Ethernet Interface on page 34

minimum-links

Syntax minimum-links *number*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces ae*x* aggregated-sonet-options],
[edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *range* aggregated-ether-options],
[edit interfaces interface-range *range* aggregated-sonet-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

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

Description For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

Options *number*—Number of links.

Range: On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled upOn EX8200 switches, the range of valid values for minimum-links number is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

Default: 1

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

Related Documentation

- Configuring Aggregated Ethernet Minimum Links on page 175
- *Configuring Aggregated SONET/SDH Interfaces*
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Junos OS Services Interfaces Library for Routing Devices*

multicast-router-interface (IGMP Snooping)

Syntax multicast-router-interface;

Hierarchy Level [edit bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping vlan *vlan-id* interface *interface-name*],
[edit protocols igmp-snooping vlan (all | *vlan-name*) interface (all | *interface-name*)]
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols vlan *vlan-id* igmp-snooping interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description Statically configure the interface as an IGMP snooping multicast-router interface—that is, an interface that faces toward a multicast router or other IGMP querier.

NOTE: If the specified interface is a trunk port, the interface becomes a multicast-routing device interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast routing device interface, even if the interface is configured as a multicast routing device interface only for IGMP snooping.

Configure an interface as a bridge interface toward other multicast routing devices.

Default The interface can either be a host-side or multicast-routing device interface.

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

Related Documentation

- *Example: Configuring IGMP Snooping on EX Series Switches*
- *Example: Configuring IGMP Snooping*
- *Configuring IGMP Snooping (CLI Procedure)*
- IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview on page 20
- *host-only-interface*
- *show igmp-snooping membership*

multi-chassis-protection

Syntax

```
multi-chassis-protection {
  peer a.b.c.d {
    interface interface-name;
    }
  }
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, you can use this statement under the physical interface level to reduce the configuration at the logical interface level if the following assumption exists:

If there are n +1 logical interfaces under **ae0**, from **ae0.0** through **ae0.n**, there will be n + 1 logical interfaces under **ge-0/0/0** as well, from **ge-0/0/0.0** through **ge-0/0/0.n**, and each **ge-0/0/0** logical interface will be a protection link for the **ae0** logical interface.

NOTE: **A bridge domain cannot have MC-AE logical interfaces which belong to different redundancy groups.**

If the Inter-Chassis Control Protocol (ICCP) connection is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer.

The remaining statements are explained separately.

Options **interface** ***interface-name***—Specify the interface: **interface** ***interface-name-fpc/pic/port***

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

Related Documentation

- Configuring Multichassis Link Aggregation on page 42
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Configuring Aggregated Ethernet Link Protection on page 146
- Example: Configuring Aggregated Ethernet Link Protection on page 148
- **peer on page 244**

non-revertive (Interfaces)

Syntax non-revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

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

Related Documentation

- link-protection on page 233
- Configuring Aggregated Ethernet Link Protection on page 146

- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

peer

Syntax

```
peer a.b.c.d {
  interface interface-name;
}
```

Hierarchy Level [edit interfaces *interface-name* multi-chassis-protection]

Release Information Statement introduced in Junos OS Release 11.1.

Description For MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces, use the **multi-chassis-protection** statement under the physical interface level to reduce the configuration at the logical interface level. If the interchassis control protocol connection (ICCP) is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement. You must also specify the peer's physical interface.

Options ***a.b.c.d***—Specify the IP address of the peer.

interface *interface-name*—Specify the peer's physical interface: interface *interface-name-fpc/pic/port*

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

Related Documentation

- Configuring Multichassis Link Aggregation on page 42
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51
- Configuring Aggregated Ethernet Link Protection on page 146
- Example: Configuring Aggregated Ethernet Link Protection on page 148
- **multi-chassis-protection on page 242**

periodic

Syntax periodic *interval*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options lacp],
[edit interfaces interface-range *name* aggregated-ether-options lacp]

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

Description For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

Options ***interval***—Interval for periodic transmission of LACP packets.

- **fast**—Transmit packets every second.
- **slow**—Transmit packets every 30 seconds.

Default: fast

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

Related Documentation

- Configuring Aggregated Ethernet LACP on page 164
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

pop-pop

Syntax pop-pop;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, and for 10-Gigabit Ethernet SFP interfaces on EX Series switches, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.

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

Related Documentation
- *Removing the Outer and Inner VLAN Tags*

pop-swap

Syntax pop-swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2, IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- *Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag*

port-priority

Syntax port-priority *priority*;

Hierarchy Level [edit interfaces *interface-name* gigether-options 802.3ad lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Define LACP port priority at the interface level.

Options *priority*—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected.

Range: 1 through 255

Default: 127

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

Related Documentation

- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*
- *Configuring Aggregated Ethernet LACP*

push-push

Syntax push-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- *Stacking Two VLAN Tags*

revertive

Syntax revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Enable the ability to switch to a better priority link (if one is available).

NOTE: By default, LACP link protection is revertive. However, you can use this statement to define a specific aggregated Ethernet interface as revertive to override a global non-revertive statement specified at the [edit chassis] hierarchy level.

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

Related Documentation

- *non-revertive (Chassis)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

swap-push

Syntax swap-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation
- *Rewriting a VLAN Tag and Adding a New Tag*

swap-swap

Syntax swap-swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and for 100-Gigabit Ethernet Type 5 PIC with CFP.

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

Related Documentation

- *Rewriting the Inner and Outer VLAN Tags*

system-id

Syntax system-id *system-id*;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp]

Release Information Statement introduced in Junos OS Release 12.2R1

Description Define the LACP system identifier at the aggregated Ethernet interface level.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

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

Related Documentation

- Configuring Aggregated Ethernet LACP on page 164

system-priority

Syntax system-priority *priority*;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global **[edit chassis]** hierarchy level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options *priority*—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

Range: 0 through 65535

Default: 127

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

Related Documentation

- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

traceoptions (LACP)

Syntax

```
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
  }
fast-hello-issu
```

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced in Junos OS Release 7.6.

Description Define tracing operations for the LACP protocol.

Default If you do not include this statement, no LACP protocol tracing operations are performed.

Options **disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, interface process tracing output is placed in the file **lacpd**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0.** When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option:

Syntax: ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

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

Related Documentation

- Tracing LACP Operations on page 170

vlan-id (VLAN ID to Be Bound to a Logical Interface)

Syntax

vlan-id *number*;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options

number—A valid VLAN identifier.

Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.

For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.

VLAN ID 0 is reserved for tagging the priority of frames.

Required Privilege Level

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

Related Documentation

- *Configuring Mixed Tagging*

vlan-tagging

Syntax vlan-tagging;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.

Description For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

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

Related Documentation

- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *Example: Configuring BGP Autodiscovery for LDP VPLS*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- Configuring Tagged Aggregated Ethernet Interfaces on page 172
- *Configuring Interfaces for VPLS Routing*
- *Enabling VLAN Tagging*
- *802.1Q VLANs Overview*
- *vlan-id*

PART 3

Administration

CHAPTER 5

Monitoring Commands

- request interface mc-ae switchover (Multichassis Link Aggregation)
- show iccp
- show interfaces (Aggregated Ethernet)
- show interfaces demux0 (Demux Interfaces)
- show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)
- show interfaces (Fast Ethernet)
- show interfaces (10-Gigabit Ethernet)
- show interfaces interface-set (Ethernet Interface Set)
- show interfaces interface-set queue
- show interfaces irb
- show l2-learning instance
- show l2-learning redundancy-groups
- show lacp interfaces
- show interfaces mac-database (Gigabit Ethernet)
- show interfaces mc-ae
- 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
- show protection-group ethernet-ring aps

- show protection-group ethernet-ring data-channel
- show protection-group ethernet-ring interface
- show protection-group ethernet-ring node-state
- show protection-group ethernet-ring statistics
- show protection-group ethernet-ring vlan

request interface mc-ae switchover (Multichassis Link Aggregation)

Syntax

request interface mc-ae switchover
<immediate> mcae-id *mcae-id*;
mcae-id *mcae-id*;

Release Information

Command introduced in Junos OS Release 13.3.

Description

Manually revert egress traffic from the active node to the designated preferred node of a multichassis aggregated Ethernet interface. You can use this command to manually switch over traffic to the preferred node when the **switchover-mode** statement for the multichassis aggregated Ethernet interface is configured as **non-revertive** at the **[edit interfaces ae*X* mc-ae**] hierarchy level.

NOTE: To run this command successfully, the **status-control** statement should be configured as **active** at the **[edit interfaces ae*X* mc-ae**] hierarchy level.

Options

immediate—(Optional) Trigger immediate switchover to the preferred node. If this option is not configured, Junos OS waits for the timer configured using the **revert-time** statement at the **[edit interfaces ae*X* mc-ae**] hierarchy level to expire before it triggers the switchover.

mcae-id *mcae-id*—Triggers switchover for the specified mc-ae interface.

Required Privilege Level

view

Related Documentation

- Configuring Multichassis Link Aggregation on page 42
- Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on page 127

List of Sample Output

request interface mc-ae switchover immediate mcae-id on page 261
request interface mc-ae switchover mcae-id on page 261

Output Fields

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

Sample Output

request interface mc-ae switchover immediate mcae-id

```
user@host >request interface mc-ae switchover immediate mcae-id 2
MCAE: Switchover Done
```

Sample Output

request interface mc-ae switchover mcae-id

```
user@host >request interface mc-ae switchover mcae-id 2
```



```
Switchover In Progress: Please check after 1 minutes,
Use “show interfaces mc-ae revertive-info” to check for the status
```


show iccp

Syntax

show iccp <brief | detail>
logical-system [*system-name* | all]

Release Information

Command introduced in Junos OS Release 10.0 for the MX Series.
Support for logical systems added in Junos OS Release 14.1.

Description

Display Interchassis Control Protocol (ICCP) information about the multichassis link aggregation group (MC-LAG) peers, including the state of the TCP connection, Bidirectional Forwarding Detection protocol, backup liveness peer status, and MCSNOOPD, LACPD, and ESWD applications.

Options

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

none—Display ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

brief—Display brief ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

detail—Display detailed ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.

Required Privilege Level

view

Related Documentation

- *Understanding Multichassis Link Aggregation*

List of Sample Output

show iccp on page 264

Output Fields

Table 5 on page 263 lists the output fields for the **show iccp** command. Output fields are listed in the approximate order in which they appear.

Table 5: show iccp

Field Name	Field Description
Redundancy Group Information for peer	Aggregated Ethernet interface name.
TCP Connection	Specifies if the TCP connection between the peers hosting the MC-LAG is up or down.
Status	Displays the state of the redundancy group: up or down
Redundancy Group ID	Displays the redundancy group identifier that is used to associate with multiple chassis that perform similar failover operations.

Table 5: show iccp *(continued)*

Field Name	Field Description
Status	Displays the state of the redundancy group: up or down
Client Application	Specifies information regarding the state of the MCSNOOPD and ESWD client applications.
Redundancy Group IDs Joined	Denotes the redundancy group unique identifier that is associated for the particular client application or process.

Sample Output

show iccp

```
user@switch> show iccp
Logical system :LS1
    Redundancy Group Information for peer 16.1.1.1
       TCP Connection       : Established
       Liveliness Detection : Up
       Redundancy Group ID          Status
         2                           Up
         12                         Up

Client Application: lacpd
Redundancy Group IDs Joined:  1
Redundancy Group IDs Joined:  2

Client Application: l2ald_iccpd_client
 Redundancy Group IDs Joined:  1
 Redundancy Group IDs Joined:  2
```


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 270
show interfaces brief (Aggregated Ethernet) on page 271
show interfaces detail (Aggregated Ethernet) on page 271
show interfaces extensive (Aggregated Ethernet) on page 272
show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 273

Output Fields

Table 6 on page 265 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 6: 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 6: 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 6: 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 6: 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 6: 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 6: 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 282
show interfaces (PPPoE over Aggregated Ethernet) on page 283
show interfaces extensive (Targeted Distribution for Aggregated Ethernet Links) on page 283
show interfaces demux0 (ACI Interface Set Configured) on page 284

Output Fields Table 7 on page 276 lists the output fields for the **show interfaces** (demux interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 7: 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 7: 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 7: 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 7: 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 7: 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 7: 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 7: 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 diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)

Syntax show interfaces diagnostics optics *interface-name*

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for PTX Series routers.

Description Display diagnostics data, warnings, and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, or 100-Gigabit Ethernet interfaces.

Options *interface-name*—Interface name. For example:

ge-*fpc/pic/port*

et-*fpc/pic/port*

xe-*fpc/pic/port*

Additional Information The transceivers are polled in 1-second intervals for diagnostics data, warnings, and alarms. The alarms do not cause the links to go down or the LEDs to change color, nor generate SNMP traps. Changes in alarm and warning status will generate system log messages.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transceiver vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a device is not working.

NOTE: Some transceivers do not support all optical diagnostics features described in the output fields.

The **show interfaces diagnostics optics** command for optical interfaces does not report the decibel (dBm) value of the received signal if the received power is zero milliwatts (0.0000 mW).

Required Privilege Level view

Related Documentation

- *Supported Network Interface Standards by Transceiver*
- *Supported Network Interface Standards by Transceiver for PTX Series Packet Transport Routers*

List of Sample Output

show interfaces diagnostics optics (DWDM and DWDM OTN) on page 298
show interfaces diagnostics optics (Bidirectional SFP) on page 298
show interfaces diagnostics optics (SFP) on page 299
show interfaces diagnostics optics (SFP) on page 299

show interfaces diagnostics optics (XFP and CFP Optics) on page 300
show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP) on page 301
show interfaces diagnostics optics for 40-Gigabit Ethernet on page 302

Output Fields Table 8 on page 286 lists the output fields for the **show interfaces diagnostics optics** command for DWDM and DWDM OTN PICs. Output fields are listed in the approximate order in which they appear.

Table 8: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the **LsPOWMON** pin in hardware.
Receiver signal average optical power	Average received optical power, in mW and dBm. This indicator is a software equivalent to the **RxPOWMON** pin in hardware. Average optical power is vendor-specific.
Laser end-of-life alarm	Laser end-of-life alarm: **On** or **Off**.
Laser wavelength alarm	Laser wavelength alarm: **On** or **Off**.
Laser bias current alarm	Laser bias current alarm: **On** or **Off**.
Laser temperature alarm	Laser temperature alarm: **On** or **Off**.
Laser power alarm	Laser power alarm: **On** or **Off**.
Modulator temperature alarm	Modulator temperature alarm: **On** or **Off**. Transceivers from some vendors do not support this field.
Modulator bias alarm	Modulator bias alarm: **On** or **Off**.
Tx multiplexer FIFO error alarm	Transmit multiplexer first in, first out (FIFO) error alarm: **On** or **Off**.
Tx loss of PLL lock alarm	Transmit loss of phase-locked loop (PLL) lock alarm: **On** or **Off**.
Rx loss of average optical power alarm	Receive loss of average optical power alarm: **On** or **Off**.

Table 8: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs *(continued)*

Field Name	Field Description
Rx loss of AC power alarm	Receive loss of AC power alarm: **On** or **Off**. Transceivers from some vendors do not support this field.
Rx loss of PLL lock alarm	Receive loss of phase-locked loop (PLL) lock alarm: **On** or **Off**.

Table 9 on page 287 lists the output fields for the **show interfaces diagnostics optics** command when the router is operating with bidirectional SFP optics. Output fields are listed in the approximate order in which they appear.

Table 9: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).
Module temperature	Temperature of the optics module, in Celsius and Fahrenheit.
Module voltage	Internally measured module voltage.
Receiver signal average optical power	Average received optical power, in mW and dBm.
Laser bias current high alarm	Laser bias power setting high alarm. Displays **on** or **off**.
Laser bias current low alarm	Laser bias power setting low alarm. Displays **on** or **off**.
Laser bias current high warning	Laser bias power setting high warning. Displays **on** or **off**.
Laser bias current low warning	Laser bias power setting low warning. Displays **on** or **off**.
Laser output power high alarm	Laser output power high alarm. Displays **on** or **off**.
Laser output power low alarm	Laser output power low alarm. Displays **on** or **off**.

Table 9: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
Laser output power high warning	Laser output power high warning. Displays **on** or **off**.
Laser output power low warning	Laser output power low warning. Displays **on** or **off**.
Module temperature high alarm	Module temperature high alarm. Displays **on** or **off**.
Module temperature low alarm	Module temperature low alarm. Displays **on** or **off**.
Module temperature high warning	Module temperature high warning. Displays **on** or **off**.
Module temperature low warning	Module temperature low warning. Displays **on** or **off**.
Module voltage high alarm	Module voltage high alarm. Displays **on** or **off**.
Module voltage low alarm	Module voltage low alarm. Displays **on** or **off**.
Module voltage high warning	Module voltage high warning. Displays **on** or **off**.
Module voltage low warning	Module voltage high warning. Displays **on** or **off**.
Laser rx power high alarm	Receive laser power high alarm. Displays **on** or **off**.
Laser rx power low alarm	Receive laser power low alarm. Displays **on** or **off**.
Laser rx power high warning	Receive laser power high warning. Displays **on** or **off**.
Laser rx power low warning	Receive laser power low warning. Displays **on** or **off**.
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm: **70.000 mA.**
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm: **0.0002 mA.**

Table 9: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning: **65.000 mA.**
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning: **0.0002 mA.**
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm: **1.0000 mW** or **0.00 dBm.**
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm: **0.0560 mW** or **-12.52 dBm.**
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning: **0.6300 mW** or **-2.01 dBm.**
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning: **0.0890 mW** or **-10.51 dBm.**
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm: **100**° C or **212**° F.
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm: **-50**° C or **-58**° F.
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning: **95** ° C or **203** ° F.
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning: **-48**° C or **-54**° F.
Module voltage high alarm threshold	Module voltage high alarm threshold: **3.700 v**.
Module voltage low alarm threshold	Module voltage low alarm threshold: **2.900 v**.
Module voltage high warning threshold	Module voltage high warning threshold: **3.7600 v**.
Module voltage low warning threshold	Module voltage low warning threshold: **3.000 v**.
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm: **1.9953 mW** or **3.00 dBm.**
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm: **0.0001 mW** or **-40.00 dBm.**

Table 9: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics *(continued)*

Field Name	Field Description
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning: **1.0000 mW** or **0.00 dBm.**
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning: **0.0010 mW** or **-30.00 dBm.**

Table 10 on page 290 lists the output fields for the **show interfaces diagnostics optics** command for SFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 10: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Measured laser bias current in uA.
Laser output power	Measured laser output power in mW.
Module temperature	Internally measured module temperature.
Module voltage	Internally measured module voltage.
Laser rx power	Measured receive optical power in mW.
Laser bias current high alarm	Laser bias current high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser bias current low alarm	Laser bias current low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power high alarm	Laser output power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power low alarm	Laser output power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp high alarm	Module temperature high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp low alarm	Module temperature low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power high alarm	Laser receive power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power low alarm	Laser receive power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.

Table 10: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(continued)*

Field Name	Field Description
Laser bias current high warning	Laser bias current high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current low warning	Laser bias current low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power high warning	Laser output power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power low warning	Laser output power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature high warning	Module temperature high warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature low warning	Module temperature low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power high warning	Laser receive power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power low warning	Laser receive power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current high alarm threshold	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current low alarm threshold	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current high warning threshold	Laser bias current high warning threshold. Warning ranges are vendor-specific.
Laser bias current low warning threshold	Laser bias current low warning threshold. Warning ranges are vendor-specific.
Laser output power high alarm threshold	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power low alarm threshold	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power high warning threshold	Laser output power high warning threshold. Warning ranges are vendor-specific.
Laser output power low warning threshold	Laser output power low warning threshold. Warning ranges are vendor-specific.

Table 10: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(continued)*

Field Name	Field Description
Module temperature high alarm threshold	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature low alarm threshold	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature high warning threshold	Module temperature high warning threshold. Warning ranges are vendor-specific.
Module temperature low warning threshold	Module temperature low warning threshold. Warning ranges are vendor-specific.
Module voltage high alarm threshold	Module voltage high alarm threshold. Alarm ranges are vendor-specific.
Module voltage low alarm threshold	Module voltage low alarm threshold. Alarm ranges are vendor-specific.
Module voltage high warning threshold	Module voltage high warning threshold. Warning ranges are vendor-specific.
Module voltage low warning threshold	Module voltage low warning threshold. Warning ranges are vendor-specific.
Laser rx power high alarm threshold	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power low alarm threshold	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power high warning threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
Laser rx power high low threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.

Table 11 on page 292 lists the output fields for the **show interfaces diagnostics optics** command for 10-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

Table 11: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Measured laser bias current in mA.

Table 11: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(continued)*

Field Name	Field Description
Laser output power	Measured laser output power in mW.
Module temperature	Internally measured module temperature.
Laser rx power	Measured receive optical power in mW.
Laser bias current high alarm	Laser bias current high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser bias current low alarm	Laser bias current low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power high alarm	Laser output power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser output power low alarm	Laser output power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp high alarm	Module temperature high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Module temp low alarm	Module temperature low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power high alarm	Laser receive power high alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser rx power low alarm	Laser receive power low alarm: **On** or **Off**. Alarm ranges are vendor-specific.
Laser bias current high warning	Laser bias current high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current low warning	Laser bias current low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power high warning	Laser output power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser output power low warning	Laser output power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature high warning	Module temperature high warning: **On** or **Off**. Warning ranges are vendor-specific.
Module temperature low warning	Module temperature low warning: **On** or **Off**. Warning ranges are vendor-specific.

Table 11: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(continued)*

Field Name	Field Description
Laser rx power high warning	Laser receive power high warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser rx power low warning	Laser receive power low warning: **On** or **Off**. Warning ranges are vendor-specific.
Laser bias current high alarm threshold	Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current low alarm threshold	Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power high alarm threshold	Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser output power low alarm threshold	Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature high alarm threshold	Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.
Module temperature low alarm threshold	Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power high alarm threshold	Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.
Laser rx power low alarm threshold	Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.
Laser bias current high warning threshold	Laser bias current high warning threshold. Warning ranges are vendor-specific.
Laser bias current low warning threshold	Laser bias current low warning threshold. Warning ranges are vendor-specific.
Laser output power high warning threshold	Laser output power high warning threshold. Warning ranges are vendor-specific.
Laser output power low warning threshold	Laser output power low warning threshold. Warning ranges are vendor-specific.
Module temperature high warning threshold	Module temperature high warning threshold. Warning ranges are vendor-specific.
Module temperature low warning threshold	Module temperature low warning threshold. Warning ranges are vendor-specific.

Table 11: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers *(continued)*

Field Name	Field Description
Laser rx power high warning threshold	Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.
Laser rx power low warning threshold	Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.

Table 12 on page 295 lists the output fields for the **show interfaces diagnostics optics** command for XFP transceivers. Output fields are listed in the approximate order in which they appear.

Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers

Field Name	Field Description
Physical interface	Name of the physical interface.
Laser bias current	Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the **LsPOWMON** pin in hardware.
Module temperature	Temperature of the XFP optics module, in Celsius and Fahrenheit.
Laser rx power	Laser received optical power, in mW and dBm.
Laser bias current high alarm	Laser bias power setting high alarm. Displays **on** or **off**.
Laser bias current low alarm	Laser bias power setting low alarm. Displays **on** or **off**.
Laser bias current high warning	Laser bias power setting high warning. Displays **on** or **off**.
Laser bias current low warning	Laser bias power setting low warning. Displays **on** or **off**.
Laser output power high alarm	Laser output power high alarm. Displays **on** or **off**.
Laser output power low alarm	Laser output power low alarm. Displays **on** or **off**.
Laser output power high warning	Laser output power high warning. Displays **on** or **off**.

Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(continued)*

Field Name	Field Description
Laser output power low warning	Laser output power low warning. Displays **on** or **off**.
Module temperature high alarm	Module temperature high alarm. Displays **on** or **off**.
Module temperature low alarm	Module temperature low alarm. Displays **on** or **off**.
Module temperature high warning	Module temperature high warning. Displays **on** or **off**.
Module temperature low warning	Module temperature low warning. Displays **on** or **off**.
Laser rx power high alarm	Receive laser power high alarm. Displays **on** or **off**.
Laser rx power low alarm	Receive laser power low alarm. Displays **on** or **off**.
Laser rx power high warning	Receive laser power high warning. Displays **on** or **off**.
Laser rx power low warning	Receive laser power low warning. Displays **on** or **off**.
Module not ready alarm	Module not ready alarm. When **on**, indicates the module has an operational fault. Displays **on** or **off**.
Module power down alarm	Module power down alarm. When **on**, module is in a limited power mode, low for normal operation. Displays on or off.
Tx data not ready alarm	Any condition leading to invalid data on the transmit path. Displays **on** or **off**.
Tx not ready alarm	Any condition leading to invalid data on the transmit path. Displays **on** or **off**.
Tx laser fault alarm	Laser fault condition. Displays **on** or **off**.
Tx CDR loss of lock alarm	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays **on** or **off**.
Rx not ready alarm	Any condition leading to invalid data on the receive path. Displays **on** or **off**.
Rx loss of signal alarm	Receive Loss of Signal alarm. When **on**, indicates insufficient optical input power to the module. Displays **on** or **off**.
Rx CDR loss of lock alarm	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays **on** or **off**.

Table 12: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers *(continued)*

Field Name	Field Description
Laser bias current high alarm threshold	Vendor-specified threshold for the laser bias current high alarm: **130.000 mA.**
Laser bias current low alarm threshold	Vendor-specified threshold for the laser bias current low alarm: **10.000 mA.**
Laser bias current high warning threshold	Vendor-specified threshold for the laser bias current high warning: **120.000 mA.**
Laser bias current low warning threshold	Vendor-specified threshold for the laser bias current low warning: **12.000 mA.**
Laser output power high alarm threshold	Vendor-specified threshold for the laser output power high alarm: **0.8910 mW** or **-0.50 dBm.**
Laser output power low alarm threshold	Vendor-specified threshold for the laser output power low alarm: **0.2230 mW** or **-6.52 dBm.**
Laser output power high warning threshold	Vendor-specified threshold for the laser output power high warning: **0.7940 mW** or **-100 dBm.**
Laser output power low warning threshold	Vendor-specified threshold for the laser output power low warning: **0.2510 mW** or **-600 dBm.**
Module temperature high alarm threshold	Vendor-specified threshold for the module temperature high alarm: **90**° C or **194**° F.
Module temperature low alarm threshold	Vendor-specified threshold for the module temperature low alarm: **-5**° C or **23**° F.
Module temperature high warning threshold	Vendor-specified threshold for the module temperature high warning: **85** ° C or **185** ° F.
Module temperature low warning threshold	Vendor-specified threshold for the module temperature low warning: **0**° C or **32**° F.
Laser rx power high alarm threshold	Vendor-specified threshold for the laser Rx power high alarm: **1.2589 mW** or **1.00 dBm.**
Laser rx power low alarm threshold	Vendor-specified threshold for the laser Rx power low alarm: **0.0323 mW** or **-14.91 dBm.**
Laser rx power high warning threshold	Vendor-specified threshold for the laser Rx power high warning: **1.1220 mW** or **0.50 dBm.**
Laser rx power low warning threshold	Vendor-specified threshold for the laser Rx power low warning: **0.0363 mW** or **-14.40 dBm.**

Sample Output

show interfaces diagnostics optics (DWDM and DWDM OTN)

```
user@host> show interfaces diagnostics optics ge-5/0/0
Physical interface: ge-5/0/0
  Laser bias current                    :  79.938 mA
  Laser output power                    :  1.592 mW / 2.02 dBm
  Receiver signal average optical power :  1.3854 mW / 1.42 dBm
  Laser end-of-life alarm               :  Off
  Laser wavelength alarm                :  Off
  Laser bias current alarm              :  Off
  Laser temperature alarm               :  Off
  Laser power alarm                     :  Off
  Modulator temperature alarm           :  Off
  Modulator bias alarm                  :  Off
  Tx multiplexer FIFO error alarm       :  Off
  Tx loss of PLL lock alarm             :  Off
  Rx loss of average optical power alarm:  Off
  Rx loss of AC power alarm             :  Off
  Rx loss of PLL lock alarm             :  Off
```

show interfaces diagnostics optics (Bidirectional SFP)

```
user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
    Laser bias current                        :  13.356 mA
    Laser output power                        :  0.2210 mW / -6.56 dBm
    Module temperature                        :  36 degrees C / 96 degrees F
    Module voltage                            :  3.2180 V
    Receiver signal average optical power     :  0.2429 mW / -6.15 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  70.000 mA
    Laser bias current low alarm threshold    :  0.002 mA
    Laser bias current high warning threshold :  65.000 mA
    Laser bias current low warning threshold  :  0.002 mA
    Laser output power high alarm threshold   :  1.0000 mW / 0.00 dBm
    Laser output power low alarm threshold    :  0.0560 mW / -12.52 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0890 mW / -10.51 dBm
    Module temperature high alarm threshold   :  100 degrees C / 212 degrees F
    Module temperature low alarm threshold    :  -50 degrees C / -58 degrees F
```

```
    Module temperature high warning threshold :  95 degrees C / 203 degrees F
    Module temperature low warning threshold  :  -48 degrees C / -54 degrees F
    Module voltage high alarm threshold       :  3.700 V
    Module voltage low alarm threshold        :  2.900 V
    Module voltage high warning threshold     :  3.600 V
    Module voltage low warning threshold      :  3.000 V
    Laser rx power high alarm threshold       :  1.9953 mW / 3.00 dBm
    Laser rx power low alarm threshold        :  0.0001 mW / -40.00 dBm
    Laser rx power high warning threshold     :  1.0000 mW / 0.00 dBm
    Laser rx power low warning threshold      :  0.0010 mW / -30.00 dBm
```

show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-0/3/0
Physical interface: ge-0/3/0
    Laser bias current                        :  23.408 mA
    Laser output power                        :  1.479 mW / 1.70 dBm
    Module temperature                        :  37 degrees C / 99 degrees F
    Laser rx power                            :  0.121 mW / -9.16 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  31.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser output power high alarm threshold   :  6.000 mW / 7.78 dBm
    Laser output power low alarm threshold    :  0.100 mW / -10.00 dBm
    Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold    :  0 degrees C / 32 degrees F
    Laser rx power high alarm threshold       :  1.000 mW / 0.00 dBm
    Laser rx power low alarm threshold        :  0.001 mW / -30.00 dBm
    Laser bias current high warning threshold :  28.000 mA
    Laser bias current low warning threshold  :  11.000 mA
    Laser output power high warning threshold :  5.000 mW / 6.99 dBm
    Laser output power low warning threshold  :  0.500 mW / -3.01 dBm
    Module temperature high warning threshold :  70 degrees C / 158 degrees F
    Module temperature low warning threshold  :  10 degrees C / 50 degrees F
    Laser rx power high warning threshold     :  0.501 mW / -3.00 dBm
    Laser rx power low warning threshold      :  0.001 mW / -28.86 dBm
```

show interfaces diagnostics optics (SFP)

```
user@host> show interfaces diagnostics optics ge-1/0/0
Physical interface: ge-1/0/0
        Laser bias current                    :  49.010 mA
        Laser output power                    :  1.263 mW / 1.01 dBm
        Module temperature                    : 17 degrees C / 62 degrees F

        Module voltage                        :  4.21 V
```

```
        Laser rx power                              :  0.060 mW / -12.21 dBm
        Laser bias current high alarm               :  Off
        Laser bias current low alarm                :  Off
        Laser output power high alarm               :  Off
        Laser output power low alarm                :  Off
        Module temperature high alarm               :  Off
        Module temperature low alarm                :  Off
        Module voltage high alarm                   :  Off
        Module voltage low alarm                    :  Off
        Laser rx power high alarm                   :  Off
        Laser rx power low alarm                    :  Off
        Laser bias current high warning             :  Off
        Laser bias current low warning              :  Off
        Laser output power high warning             :  Off
        Laser output power low warning              :  Off
        Module temperature high warning             :  Off
        Module temperature low warning              :  Off
        Module voltage high warning                 :  Off
        Module voltage low warning                  :  Off
        Laser rx power high warning                 :  Off
        Laser rx power low warning                  :  Off
        Laser bias current high alarm threshold     :  70.000 mA
        Laser bias current low alarm threshold      :  20.000 mA
        Laser bias current high warning threshold   :  65.000 mA
        Laser bias current low warning threshold    :  25.000 mA
        Laser output power high alarm threshold     :  1.4120 mW / 1.50 dBm
        Laser output power low alarm threshold      :  0.1990 mW / -7.01 dBm
        Laser output power high warning threshold   :  1.2580 mW / 1.00 dBm
        Laser output power low warning threshold    :  0.2230 mW / -6.52 dBm
        Module temperature high alarm threshold     :  78 degrees C /172 degrees F

        Module temperature low alarm threshold      :  13 degrees C / 9 degrees F
        Module temperature high warning threshold   :  75 degrees C /167 degrees F

        Module temperature low warning threshold    :  10 degrees C / 14 degrees F

        Module voltage high alarm threshold         :  5.71 V
        Module voltage low alarm threshold          :  2.05 V
        Module voltage high warning threshold       :  5.20 V
        Module voltage low warning threshold        :  3.11 V
        Laser rx power high alarm threshold         :  1.7783 mW / 2.50 dBm
        Laser rx power low alarm threshold          :  0.0100 mW / -20.00 dBm
        Laser rx power high warning threshold       :  1.5849 mW / 2.00 dBm
        Laser rx power low warning threshold        :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics (XFP and CFP Optics)

```
user@host> show interfaces diagnostics optics xe-2/1/0
Physical interface: xe-2/1/0
    Laser bias current                          :  52.060 mA
    Laser output power                          :  0.5640 mW / -2.49 dBm
    Module temperature                          :  31 degrees C / 88 degrees F
    Laser rx power                              :  0.0844 mW / -10.74 dBm
    Laser bias current high alarm               :  Off
    Laser bias current low alarm                :  Off
    Laser bias current high warning             :  Off
    Laser bias current low warning              :  Off
    Laser output power high alarm               :  Off
    Laser output power low alarm                :  Off
    Laser output power high warning             :  Off
    Laser output power low warning              :  Off
```

```
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Module not ready alarm                    :  Off
    Module power down alarm                   :  Off
    Tx data not ready alarm                   :  Off
    Tx not ready alarm                        :  Off
    Tx laser fault alarm                      :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx not ready alarm                        :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    Laser bias current high alarm threshold   :  130.000 mA
    Laser bias current low alarm threshold    :  10.000 mA
    Laser bias current high warning threshold :  120.000 mA
    Laser bias current low warning threshold  :  12.000 mA
    Laser output power high alarm threshold   :  0.8910 mW / -0.50 dBm
    Laser output power low alarm threshold    :  0.2230 mW / -6.52 dBm
    Laser output power high warning threshold :  0.7940 mW / -1.00 dBm
    Laser output power low warning threshold  :  0.2510 mW / -6.00 dBm
    Module temperature high alarm threshold   :  90 degrees C / 194 degrees F
    Module temperature low alarm threshold    :  -5 degrees C / 23 degrees F
    Module temperature high warning threshold :  85 degrees C / 185 degrees F
    Module temperature low warning threshold  :  0 degrees C / 32 degrees F
    Laser rx power high alarm threshold       :  1.2589 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0323 mW / -14.91 dBm
    Laser rx power high warning threshold     :  1.1220 mW / 0.50 dBm
    Laser rx power low warning threshold      :  0.0363 mW / -14.40 dBm
```

show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP)

```
user@host> show interfaces diagnostics optics et-2/0/23
Physical interface: et-2/0/23
    Laser bias current                        :  8.482 mA
    Laser output power                        :  0.5890 mW / -2.30 dBm
    Module temperature                        :  51 degrees C / 123 degrees F
    Module voltage                            :  3.2970 V
    Receiver signal average optical power     :  0.5574 mW / -2.54 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
```

```
    Laser rx power high warning                  :  Off
    Laser rx power low warning                   :  Off
    Laser bias current high alarm threshold      :  11.800 mA
    Laser bias current low alarm threshold       :  4.000 mA
    Laser bias current high warning threshold    :  10.800 mA
    Laser bias current low warning threshold     :  5.000 mA
    Laser output power high alarm threshold      :  0.8310 mW / -0.80 dBm
    Laser output power low alarm threshold       :  0.2510 mW / -6.00 dBm
    Laser output power high warning threshold    :  0.6600 mW / -1.80 dBm
    Laser output power low warning threshold     :  0.3160 mW / -5.00 dBm
    Module temperature high alarm threshold      :  93 degrees C / 199 degrees F
    Module temperature low alarm threshold       :  -13 degrees C / 9 degrees F
    Module temperature high warning threshold    :  88 degrees C / 190 degrees F
    Module temperature low warning threshold     :  -8 degrees C / 18 degrees F
    Module voltage high alarm threshold          :  3.700 V
    Module voltage low alarm threshold           :  2.900 V
    Module voltage high warning threshold        :  3.600 V
    Module voltage low warning threshold         :  3.000 V
    Laser rx power high alarm threshold          :  1.0000 mW / 0.00 dBm
    Laser rx power low alarm threshold           :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold        :  0.7943 mW / -1.00 dBm
    Laser rx power low warning threshold         :  0.0158 mW / -18.01 dBm
```

show interfaces diagnostics optics for 40-Gigabit Ethernet

```
user@host> show interfaces diagnostics optics et-7/1/0
Physical interface: et-7/1/0
    Module temperature                           :  34 degrees C / 94 degrees F
    Module voltage                               :  3.4720 V
    Module temperature high alarm                :  Off
    Module temperature low alarm                 :  Off
    Module temperature high warning              :  Off
    Module temperature low warning               :  Off
    Module voltage high alarm                    :  Off
    Module voltage low alarm                     :  Off
    Module voltage high warning                  :  Off
    Module voltage low warning                   :  Off
    Module not ready alarm                       :  Off
    Module low power alarm                       :  Off
    Module initialization incomplete alarm       :  Off
    Module fault alarm                           :  Off
    PLD Flash initialization fault alarm         :  Off
    Power supply fault alarm                     :  Off
    Checksum fault alarm                         :  Off
    Tx laser disabled alarm                      :  Off
    Tx loss of signal functionality alarm        :  Off
    Tx CDR loss of lock alarm                    :  Off
    Rx loss of signal alarm                      :  Off
    Rx CDR loss of lock alarm                    :  Off
    Module temperature high alarm threshold      :  80 degrees C / 176 degrees F
    Module temperature low alarm threshold       :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold    :  75 degrees C / 167 degrees F
    Module temperature low warning threshold     :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold          :  3.5990 V
    Module voltage low alarm threshold           :  3.0000 V
    Module voltage high warning threshold        :  3.5000 V
    Module voltage low warning threshold         :  3.0990 V
    Laser bias current high alarm threshold      :  100.000 mA
    Laser bias current low alarm threshold       :  10.000 mA
    Laser bias current high warning threshold    :  80.000 mA
```



```
    Laser bias current low warning threshold  :  15.000 mA
    Laser output power high alarm threshold   :  2.8180 mW / 4.50 dBm
    Laser output power low alarm threshold    :  0.2390 mW / -6.22 dBm
    Laser output power high warning threshold :  2.2380 mW / 3.50 dBm
    Laser output power low warning threshold  :  0.3010 mW / -5.21 dBm
    Laser rx power high alarm threshold       :  2.5119 mW / 4.00 dBm
    Laser rx power low alarm threshold        :  0.0316 mW / -15.00 dBm
    Laser rx power high warning threshold     :  1.9953 mW / 3.00 dBm
    Laser rx power low warning threshold      :  0.0631 mW / -12.00 dBm
    Laser temperature high alarm threshold    :  80 degrees C / 176 degrees F
    Laser temperature low alarm threshold     :  -10 degrees C / 14 degrees F
    Laser temperature high warning threshold  :  75 degrees C / 167 degrees F
    Laser temperature low warning threshold   :  -5 degrees C / 23 degrees F
  Lane 0
    Laser bias current                        :  27.829 mA
    Laser output power                        :  0.851 mW / -0.70 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.894 mW / -0.49 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
    Laser receiver power low warning          :  Off
    Tx loss of signal functionality alarm     :  Off
    Tx CDR loss of lock alarm                 :  Off
    Rx loss of signal alarm                   :  Off
    Rx CDR loss of lock alarm                 :  Off
    APD supply fault alarm                    :  Off
    TEC fault alarm                           :  Off
    Wavelength unlocked alarm                 :  Off
  Lane 1
    Laser bias current                        :  35.374 mA
    Laser output power                        :  0.896 mW / -0.48 dBm
    Laser temperature                         :  34 degrees C / 94 degrees F
    Laser receiver power                      :  0.707 mW / -1.50 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Laser temperature high alarm              :  Off
    Laser temperature low alarm               :  Off
    Laser temperature high warning            :  Off
    Laser temperature low warning             :  Off
    Laser receiver power high alarm           :  Off
    Laser receiver power low alarm            :  Off
    Laser receiver power high warning         :  Off
```

```
      Laser receiver power low warning         :  Off
      Tx loss of signal functionality alarm    :  Off
      Tx CDR loss of lock alarm                :  Off
      Rx loss of signal alarm                  :  Off
      Rx CDR loss of lock alarm                :  Off
      APD supply fault alarm                   :  Off
      TEC fault alarm                          :  Off
      Wavelength unlocked alarm                :  Off
    Lane 2
      Laser bias current                       :  29.173 mA
      Laser output power                       :  0.890 mW / -0.51 dBm
      Laser temperature                        :  34 degrees C / 94 degrees F
      Laser receiver power                     :  0.704 mW / -1.52 dBm
      Laser bias current high alarm            :  Off
      Laser bias current low alarm             :  Off
      Laser bias current high warning          :  Off
      Laser bias current low warning           :  Off
      Laser output power high alarm            :  Off
      Laser output power low alarm             :  Off
      Laser output power high warning          :  Off
      Laser output power low warning           :  Off
      Laser temperature high alarm             :  Off
      Laser temperature low alarm              :  Off
      Laser temperature high warning           :  Off
      Laser temperature low warning            :  Off
      Laser receiver power high alarm          :  Off
      Laser receiver power low alarm           :  Off
      Laser receiver power high warning        :  Off
      Laser receiver power low warning         :  Off
      Tx loss of signal functionality alarm    :  Off
      Tx CDR loss of lock alarm                :  Off
      Rx loss of signal alarm                  :  Off
      Rx CDR loss of lock alarm                :  Off
      APD supply fault alarm                   :  Off
      TEC fault alarm                          :  Off
      Wavelength unlocked alarm                :  Off
    Lane 3
      Laser bias current                       :  36.164 mA
      Laser output power                       :  0.899 mW / -0.46 dBm
      Laser temperature                        :  34 degrees C / 94 degrees F
      Laser receiver power                     :  0.892 mW / -0.50 dBm
      Laser bias current high alarm            :  Off
      Laser bias current low alarm             :  Off
      Laser bias current high warning          :  Off
      Laser bias current low warning           :  Off
      Laser output power high alarm            :  Off
      Laser output power low alarm             :  Off
      Laser output power high warning          :  Off
      Laser output power low warning           :  Off
      Laser temperature high alarm             :  Off
      Laser temperature low alarm              :  Off
      Laser temperature high warning           :  Off
      Laser temperature low warning            :  Off
      Laser receiver power high alarm          :  Off
      Laser receiver power low alarm           :  Off
      Laser receiver power high warning        :  Off
      Laser receiver power low warning         :  Off
      Tx loss of signal functionality alarm    :  Off
      Tx CDR loss of lock alarm                :  Off
      Rx loss of signal alarm                  :  Off
      Rx CDR loss of lock alarm                :  Off
```



```
APD supply fault alarm                    :  Off
TEC fault alarm                           :  Off
Wavelength unlocked alarm                 :  Off
```

show interfaces (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 319
show interfaces brief (Fast Ethernet) on page 320
show interfaces detail (Fast Ethernet) on page 320
show interfaces extensive (Fast Ethernet) on page 320

Output Fields Table 13 on page 306 lists the output fields for the **show interfaces Fast Ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 13: 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 13: 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 13: 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 13: 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 13: 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 13: 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 13: 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 13: 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 13: show interfaces Fast Ethernet Output Fields *(continued)*

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

Table 13: 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 13: 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 13: 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 13: 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 13: 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

Output Fields

See Table 14 on page 323 for the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Table 14: 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 14: 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 14: 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 14 on page 323.	**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 14: 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 14: 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 14: 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 14: 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 15 on page 337 • **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 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 14: 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) • SE**S-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 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 14: 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 14: 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 14: 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 14: 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 14: 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 15 on page 337 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 15 on page 337, 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 15: 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 348
show interfaces interface-set detail on page 348
show interfaces interface-set (ACI Interface Set) on page 349

Output Fields Table 16 on page 347 describes the information for the **show interfaces interface-set** command.

Table 16: 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 16: 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 interfaces interface-set queue

Syntax

show interfaces interface-set queue *interface-set-name*
<aggregate | remaining-traffic>
<forwarding-class *class-name*>

Release Information

Command introduced in Junos OS Release 8.5.

Description

Display information about the gigabit or 10-Gigabit Ethernet interface set queue. Supported in MX Series routers with enhanced queuing DPCs.

Options

interface-set-name—(Optional) Display information about the specified gigabit or 10-Gigabit Ethernet interface set. Wildcard values can be used in the interface set name.

aggregate—(Optional) Display the aggregated queuing statistics of all member logical interfaces for interface sets that have traffic-control profiles configured.

both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics.

egress—(Optional) Display egress queue statistics.

forwarding-class *class-name*—(Optional) Display queuing statistics for the specified forwarding class.

ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics.

remaining-traffic—(Optional) Display the queuing statistics of all member logical interfaces for interface sets that do not have traffic-control profiles configured.

Required Privilege Level

view

List of Sample Output

show interfaces interface-set queue (Gigabit Ethernet) on page 351
show interfaces interface-set queue both-ingress-egress (Enhanced DPC) on page 352
show interfaces interface-set queue egress (Enhanced DPC) on page 354
show interfaces interface-set queue forwarding-class (Gigabit Ethernet) on page 355
show interfaces interface-set queue (Enhanced DPC) on page 356
show interfaces interface-set queue remaining-traffic (Gigabit Ethernet) on page 356

Output Fields

Table 17 on page 350 describes the information for the **show interfaces interface-set queue** command.

Table 17: Ethernet show interfaces interface-set queue Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Interface set	Name of the interface set.	All levels
Interface set index	Index number of the interface set.	All levels

Table 17: Ethernet show interfaces interface-set queue Output Fields *(continued)*

Field Name	Field Description	Level of Output
Forwarding classes supported	Total number of forwarding classes supported on the specified interface set.	All levels
Forwarding classes in use	Total number of forwarding classes used on the specified interface set.	All levels
Egress queues supported	Total number of egress queues supported on the specified interface set.	All levels
Egress queues in use	Total number of egress queues used on the specified interface set.	All levels
Ingress queues supported	Total number of ingress queues supported on the specified interface set.	All levels
Ingress queues in use	Total number of ingress queues used on the specified interface set.	All levels
Queue	Egress or ingress queue number for the statistics being displayed.	All levels
Forwarding classes	Forwarding class name for the statistics being displayed.	All levels
Queued	**Packet** and **Byte** statistics for the specified queue. • **Packets**—Number of packets queued and input rate in packets per second. • **Bytes**—Number of bytes queued and input rate in bytes per second.	All levels
Transmitted	**Packet** and **Byte** statistics for the specified forwarding class. • **Packets**—Number of packets transmitted and transmit rate in packets per second. • **Bytes**—Number of bytes transmitted and transmit rate in bytes per second. • **Tail-dropped packets**—Number of packets tail dropped. • **RED-dropped packets**—Number of RED-dropped packets for the **low**, **medium-low**, **medium-high**, and **high** loss priorities. • **RED-dropped bytes**—Number of RED-dropped bytes for the **low**, **medium-low**, **medium-high**, and **high** loss priorities.	All levels

Sample Output

show interfaces interface-set queue (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :               3998482                     1 pps
```

```
    Bytes                :            271896884                   688 bps
  Transmitted:
    Packets              :              1077474                     1 pps
    Bytes                :             73268340                   688 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :              2921008                     0 pps
     Low                 :              2921008                     0 pps
     Medium-low          :                    0                     0 pps
     Medium-high         :                    0                     0 pps
     High                :                    0                     0 pps
    RED-dropped bytes    :            198628544                     0 bps
     Low                 :            198628544                     0 bps
     Medium-low          :                    0                     0 bps
     Medium-high         :                    0                     0 bps
     High                :                    0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
...
```

show interfaces interface-set queue both-ingress-egress (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 both-ingress-egress
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            185968478                473161 pps
    Bytes                :          10042313520             204441336 bps
  Transmitted:
    Packets              :              5441673                 13780 pps
    Bytes                :            293850342               5952960 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :            180526772                459372 pps
    RED-dropped bytes    :           9748446282             198451512 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :            522021472                473602 pps
    Bytes                :          28190332480             204599944 bps
  Transmitted:
    Packets              :              5791772                  4055 pps
    Bytes                :            312755688               1751976 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :            516227139                469546 pps
    RED-dropped bytes    :          27876265560             202843872 bps
Queue: 3, Forwarding classes: network-control
  Queued:
```



```
    Packets           :                0                    0 pps
    Bytes             :                0                    0 bps
  Transmitted:
    Packets           :                0                    0 pps
    Bytes             :                0                    0 bps
    Tail-dropped packets :             0                    0 pps
    RED-dropped packets  :             0                    0 pps
    RED-dropped bytes    :             0                    0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets           :          5417304                13797 pps
    Bytes             :        368429508              7506096 bps
  Transmitted:
    Packets           :          5014996                12769 pps
    Bytes             :        341019728              6946560 bps
    Tail-dropped packets :             0                    0 pps
    RED-dropped packets  :        402189                 1028 pps
     Low              :           402189                 1028 pps
     Medium-low       :                0                    0 pps
     Medium-high      :                0                    0 pps
     High             :                0                    0 pps
    RED-dropped bytes    :      27348852               559536 bps
     Low              :         27348852               559536 bps
     Medium-low       :                0                    0 bps
     Medium-high      :                0                    0 bps
     High             :                0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets           :                0                    0 pps
    Bytes             :                0                    0 bps
  Transmitted:
    Packets           :                0                    0 pps
    Bytes             :                0                    0 bps
    Tail-dropped packets :             0                    0 pps
    RED-dropped packets  :             0                    0 pps
     Low              :                0                    0 pps
     Medium-low       :                0                    0 pps
     Medium-high      :                0                    0 pps
     High             :                0                    0 pps
    RED-dropped bytes    :             0                    0 bps
     Low              :                0                    0 bps
     Medium-low       :                0                    0 bps
     Medium-high      :                0                    0 bps
     High             :                0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets           :          5770534                 3963 pps
    Bytes             :        396943252              2156144 bps
  Transmitted:
    Packets           :          3945152                 1457 pps
    Bytes             :        268270336               792608 bps
    Tail-dropped packets :             0                    0 pps
    RED-dropped packets  :       1815141                 2506 pps
     Low              :          1815141                 2506 pps
     Medium-low       :                0                    0 pps
     Medium-high      :                0                    0 pps
     High             :                0                    0 pps
    RED-dropped bytes    :     123429524              1363536 bps
     Low              :        123429524              1363536 bps
```

```
      Medium-low          :                    0                     0 bps
      Medium-high         :                    0                     0 bps
      High                :                    0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
  Transmitted:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
    Tail-dropped packets  :                    0                     0 pps
    RED-dropped packets   :                    0                     0 pps
      Low                 :                    0                     0 pps
      Medium-low          :                    0                     0 pps
      Medium-high         :                    0                     0 pps
      High                :                    0                     0 pps
    RED-dropped bytes     :                    0                     0 bps
      Low                 :                    0                     0 bps
      Medium-low          :                    0                     0 bps
      Medium-high         :                    0                     0 bps
      High                :                    0                     0 bps
```

show interfaces interface-set queue egress (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 egress
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets               :              3958253                 13822 pps
    Bytes                 :            269217592               7519712 bps
  Transmitted:
    Packets               :              3665035                 12729 pps
    Bytes                 :            249222380               6924848 bps
    Tail-dropped packets  :                    0                     0 pps
    RED-dropped packets   :               293091                  1093 pps
      Low                 :               293091                  1093 pps
      Medium-low          :                    0                     0 pps
      Medium-high         :                    0                     0 pps
      High                :                    0                     0 pps
    RED-dropped bytes     :             19930188                594864 bps
      Low                 :             19930188                594864 bps
      Medium-low          :                    0                     0 bps
      Medium-high         :                    0                     0 bps
      High                :                    0                     0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
  Transmitted:
    Packets               :                    0                     0 pps
    Bytes                 :                    0                     0 bps
    Tail-dropped packets  :                    0                     0 pps
    RED-dropped packets   :                    0                     0 pps
      Low                 :                    0                     0 pps
      Medium-low          :                    0                     0 pps
      Medium-high         :                    0                     0 pps
      High                :                    0                     0 pps
    RED-dropped bytes     :                    0                     0 bps
```

```
      Low                 :                  0                  0 bps
      Medium-low          :                  0                  0 bps
      Medium-high         :                  0                  0 bps
      High                :                  0                  0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets               :            5350989               3904 pps
    Bytes                 :          368412924            2124048 bps
  Transmitted:
    Packets               :            3790469               1465 pps
    Bytes                 :          257751892             796960 bps
    Tail-dropped packets  :                  0                  0 pps
    RED-dropped packets   :            1550282               2439 pps
     Low                  :            1550282               2439 pps
     Medium-low           :                  0                  0 pps
     Medium-high          :                  0                  0 pps
     High                 :                  0                  0 pps
    RED-dropped bytes     :          105419176            1327088 bps
     Low                  :          105419176            1327088 bps
     Medium-low           :                  0                  0 bps
     Medium-high          :                  0                  0 bps
     High                 :                  0                  0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets               :                  0                  0 pps
    Bytes                 :                  0                  0 bps
  Transmitted:
    Packets               :                  0                  0 pps
    Bytes                 :                  0                  0 bps
    Tail-dropped packets  :                  0                  0 pps
    RED-dropped packets   :                  0                  0 pps
     Low                  :                  0                  0 pps
     Medium-low           :                  0                  0 pps
     Medium-high          :                  0                  0 pps
     High                 :                  0                  0 pps
    RED-dropped bytes     :                  0                  0 bps
     Low                  :                  0                  0 bps
     Medium-low           :                  0                  0 bps
     Medium-high          :                  0                  0 bps
     High                 :                  0                  0 bps
```

show interfaces interface-set queue forwarding-class (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 forwarding-class best-effort
Interface set: ge-2/2/0-0
  Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets               :          101857694            1420083 pps
    Bytes                 :         6927234456          772532320 bps
  Transmitted:
    Packets               :            3984693              55500 pps
    Bytes                 :          270959592           30192512 bps
    Tail-dropped packets  :                  0                  0 pps
    RED-dropped packets   :           97870952            1364583 pps
     Low                  :           97870952            1364583 pps
     Medium-low           :                  0                  0 pps
     Medium-high          :                  0                  0 pps
     High                 :                  0                  0 pps
```

```
    RED-dropped bytes    :           6655225776             742339808 bps
     Low                 :           6655225776             742339808 bps
     Medium-low          :                    0                     0 bps
     Medium-high         :                    0                     0 bps
     High                :                    0                     0 bps
```

show interfaces interface-set queue (Enhanced DPC)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 ingress
Interface set: foo
  Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            149036817                473711 pps
    Bytes                :           8048003934             204642936 bps
  Transmitted:
    Packets              :              4360749                 13891 pps
    Bytes                :            235480446               6000912 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :            144676035                459820 pps
    RED-dropped bytes    :           7812506592             198642024 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :            485089207                473605 pps
    Bytes                :          26195987476             204597576 bps
  Transmitted:
    Packets              :              5480799                  3959 pps
    Bytes                :            295963146               1710504 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :            479605853                469646 pps
    RED-dropped bytes    :          25898716170             202887072 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
  Transmitted:
    Packets              :                    0                     0 pps
    Bytes                :                    0                     0 bps
    Tail-dropped packets :                    0                     0 pps
    RED-dropped packets  :                    0                     0 pps
    RED-dropped bytes    :                    0                     0 bps
```

show interfaces interface-set queue remaining-traffic (Gigabit Ethernet)

```
user@host> show interfaces interface-set queue ge-2/2/0-0 remaining-traffic
Interface set: ge-2/2/0-0
  Interface set index: 12
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
```



```
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :               2201552                     0 pps
    Bytes                :             149705536                     0 bps
  Transmitted:
    Packets              :                609765                     0 pps
    Bytes                :              41464020                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :               1591787                     0 pps
     Low                 :               1591787                     0 pps
     Medium-low          :                     0                     0 pps
     Medium-high         :                     0                     0 pps
     High                :                     0                     0 pps
    RED-dropped bytes    :             108241516                     0 bps
     Low                 :             108241516                     0 bps
     Medium-low          :                     0                     0 bps
     Medium-high         :                     0                     0 bps
     High                :                     0                     0 bps
```


show interfaces irb

Syntax

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

Release Information Command introduced in Junos OS Release 8.4.

Description Display integrated routing and bridging interfaces information.

Options

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

descriptions—(Optional) Display interface description strings.

mac—Display hardware MAC address

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

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

statistics—(Optional) Display static interface statistics.

Additional Information Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another bridging domain that has a Layer 3 protocol configured.

Required Privilege Level view

List of Sample Output

show interfaces irb extensive on page 362
show interfaces irb snmp-index on page 363

Output Fields Table 18 on page 358 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 18: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.	All levels
Proto	Protocol configured on the interface.	**terse**

Table 18: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Physical interface type.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**detail extensive brief** none
MTU	MTU size on the physical interface.	**detail extensive brief** none
Clocking	Reference clock source: **Internal** or **External**. Always unspecified on IRB interfaces.	**detail extensive brief**
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	**detail extensive brief**
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.	**detail extensive brief** none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	**detail extensive brief** none
Link type	Physical interface link type: **full duplex** or **half duplex**.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.	**detail extensive** none
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Alternate link address	Backup address of the link.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 18: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Logical Interface		

Table 18: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface (which reflects its initialization sequence).	**detail extensive** none
SNMP ifIndex	SNMP interface index number of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	**detail extensive**
Encapsulation	Encapsulation on the logical interface.	**detail extensive**
Bandwidth	Speed at which the interface is running.	**detail extensive**
Routing Instance	Routing instance IRB is configured under.	**detail extensive**
Bridging Domain	Bridging domain IRB is participating in.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router.	**detail extensive**
Protocol	Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive**
MTU	Maximum transmission unit size on the logical interface.	**detail extensive**
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none

Table 18: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under *Common Output Fields Description*.	**detail extensive**
Policer	The policer that is to be evaluated when packets are received or transmitted on the interface.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.	**detail extensive**

Sample Output

show interfaces irb extensive

```
user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
```

```
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
       IPv6 transit statistics:
        Input  bytes  :                  0
        Output bytes  :                  0
        Input  packets:                  0
        Output packets:                  0
      Local statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
      Transit statistics:
       Input  bytes  :                   0                    0 bps
       Output bytes  :                   0                    0 bps
       Input  packets:                   0                    0 pps
       Output packets:                   0                    0 pps
       IPv6 transit statistics:
        Input  bytes  :                  0
        Output bytes  :                  0
        Input  packets:                  0
        Output packets:                  0
      Protocol inet, MTU: 1500, Generation: 154, Route table: 0
        Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
          Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
          Generation: 155
      Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
        Flags: Is-Primary
        Policer: Input: __default_arp_policer
```

show interfaces irb snmp-index

```
user@host> show interfaces snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
    Protocol multiservice, MTU: 1500
      Flags: Is-Primary
```

show l2-learning instance

Syntax **show l2-learning instance**

Release Information (MX Series routers only) Command introduced in Junos OS Release 8.4.

Description Display Layer 2 learning properties for all the configured routing instances.

Options This command has no options.

Required Privilege Level view

List of Sample Output show l2-learning instance on page 364

Output Fields Table 19 on page 364 describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

Table 19: show l2-learning instance Output Fields

Field Name	Field Description
Routing Instance	Name of routing instance.
Bridging Domain	Name of bridging domain. On MX Series routers you can use the **show l2-learning instance <extensive>** command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID.
Index	Number associated with the routing instance or bridging domain.
Logical System	Name of logical system or **Default** if no logical system is configured.
Routing instance flags	Status of Layer 2 learning properties for each routing instance: • **DL**—MAC learning is disabled. • **SE**—MAC accounting is enabled. • **AD**—Packets are dropped after MAC address limit is reached. • **LH**—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.
MAC limit	Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.

Sample Output

show l2-learning instance

```
user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
```



```
                         AD -packet action drop, LH -mac limit hit)

Routing              Bridging     Index   Logical       Routing     MAC
Instance             Domain               System        flags       limit
__juniper_private1__                  1   Default                    5000
vs1                  vlan100          3   Default                    5120
vs1                  vlan200          4   Default                    5120
```


show l2-learning redundancy-groups

Syntax

show l2-learning redundancy-groups
logical-system [*system-name* | all]
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs

Release Information

Command introduced in Junos OS Release 13.2.
Support for logical systems added in Junos OS Release 14.1.

Description

(MX Series routers only) Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregate Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or logical systems on a router. Note that the Redundancy Group ID is inherited by the bridging domain from member AE interfaces.

Options

logical-system [*system-name* | all]—(Optional) Display information for a specified logical system or all systems.

redundancy-group-id—(Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.

arp-statistics—(Optional) Count of ARP packets sent and received by the two MC-AE nodes.

nd-statistics—(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.

remote-macs—(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

Required Privilege Level

view

Related Documentation

- Configuring Multichassis Link Aggregation on page 42
- **show interfaces mc-ae on page 378**
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 51

List of Sample Output

show l2-learning redundancy-groups arp-statistics on page 368
show l2-learning redundancy-groups nd-statistics on page 368
show l2-learning redundancy-groups remote-macs on page 368
show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems) on page 369
show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems) on page 369
show l2-learning redundancy-groups group-id on page 369

show l2-learning redundancy-groups logical-system on page 369

Output Fields Output fields are listed in the approximate order in which they appear.

Table 20: show l2-learning redundancy-groups arp-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ARP Statistics Group ID	ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ARP Rx Count From Line	Total number of ARPs received from the Line.
ARP Tx Count To Peer	Total number of ARPs sent to the peer.
ARP Rx Count From Peer	Total number of ARPs received from the peer.
ARP Drop Count received from line	Total number of ARPs sent by the peer that were received.
ARP Drop Count received from peer	Total number of ARPs sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 21: show l2-learning redundancy-groups nd-statistics Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
MCLAG ND Statistics Group ID	Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MC-LAG) instance.
ND Rx Count From Line	Total number of Neighbor Discovery packets received from the Line.
ND Tx Count To Peer	Total number of Neighbor Discovery packets sent to the peer.
ND Rx Count From Peer	Total number of Neighbor Discovery packets received from the peer.
ND Drop Count received from line	Total number of Neighbor Discovery packets sent by the peer that were received.
ND Drop Count received from peer	Total number of Neighbor Discovery packets sent by the peer that were dropped
Service-id	Service ID (configured at the routing instance level).

Table 22: show l2-learning redundancy-groups remote-macs Output Fields

Field Name	Field Description
Redundancy Group ID	Redundancy Group to which the following details apply.
Peer-Addr	IP address of the remote peer.
VLAN	Virtual LAN identifier associated with the redundancy group.
MAC	Hardware media access control address associated with the redundancy group.
MCAE-ID	ID number of the MC-AE used by the redundancy group.
Flags	Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected.
Status	Installation state: Installed or Not Installed.

Sample Output

show l2-learning redundancy-groups arp-statistics

```
user@host> show l2-learning redundancy-groups arp-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                : 1
ARP Rx Count From Line                  : 52
ARP Tx Count To Peer                    : 15
ARP Rx Count From Peer                  : 39
ARP Install Count                       : 34
ARP Drop Count received from line       : 37
ARP Drop Count received from peer       : 5
```

show l2-learning redundancy-groups nd-statistics

```
user@host> show l2-learning redundancy-groups nd-statistics
Logical System : default
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                : 1
ND Rx Count From Line                  : 52
ND Tx Count To Peer                    : 15
ND Rx Count From Peer                  : 39
ND Install Count                       : 34
ND Drop Count received from line       : 37
ND Drop Count received from peer       : 5
```

show l2-learning redundancy-groups remote-macs

```
user@host> show l2-learning redundancy-groups <redundancy-group-id> remote-macs
  Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

Service-id  Peer-Addr    VLAN        MAC            MCAE-ID  Subunit  Opcode
```

```
Flags   Status
10         1.1.1.2      100   64:87:88:6a:df:f0      1        0       1
0     Installed
```

show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems)

user@host> **show l2-learning redundancy-groups logical-system LS1 arp-statistics**

```
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID                                : 1
ARP Rx Count From Line                  : 52
ARP Tx Count To Peer                    : 15
ARP Rx Count From Peer                  : 39
ARP Install Count                       : 34
ARP Drop Count received from line       : 37
ARP Drop Count received from peer       : 5
```

show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems)

user@host> **show l2-learning redundancy-groups logical-system LS1 nd-statistics**

```
 Redundancy Group ID : 1     Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID                                : 1
ND Rx Count From Line                  : 52
ND Tx Count To Peer                    : 15
ND Rx Count From Peer                  : 39
ND Install Count                       : 34
ND Drop Count received from line       : 37
ND Drop Count received from peer       : 5
```

show l2-learning redundancy-groups group-id

user@host> **show l2-learning redundancy-groups 1**

```
Redundancy Group ID : 1       Flags : Local Connect,Remote Connect
```

show l2-learning redundancy-groups logical-system

user@host> **show l2-learning redundancy-groups logical-system ls1**

```
Redundancy Group ID : 2       Flags : Local Connect,Remote Connect
```

show lacp interfaces

Syntax

show lacp interfaces
<*interface-name*>

Release Information

Command introduced in Junos OS Release 7.6.

Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

Options

none—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*number***
- Fast Ethernet—**fe-*fpc/pic/port***
- Gigabit Ethernet—**ge-*fpc/pic/port***

NOTE: The **show lacp interfaces** command returns the following error message if your system is not configured in either active or passive LACP mode:

"Warning: lacp subsystem not running – not needed by configuration"

Required Privilege Level

view

List of Sample Output

show lacp interfaces (Aggregated Ethernet) on page 373
show lacp interfaces (Gigabit Ethernet) on page 373

Output Fields

Table 5 on page 263 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 23: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Table 23: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP State	LACP state information for each aggregated interface:
	• **Role**—Role played by the interface. It can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **Exp**—Expired state. **Yes** indicates the actor or partner is in an expired state. **No** indicates the actor or partner is not in an expired state.
	• **Def**—Default. **Yes** indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. **No** indicates the operational partner information in use has been received in an LACP PDU.
	• **Dist**—Distribution of outgoing frames. **No** indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is **Yes**.
	• **Col**—Collection of incoming frames. **Yes** indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is **No**.
	• **Syn**—Synchronization. If the value is **Yes**, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is **No**, the link is not synchronized. It is currently not in the right aggregation.
	• **Aggr**—Ability of aggregation port to aggregate (**Yes**) or to operate only as an individual link (**No**).
	• **Timeout**—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (**Long Timeout** or **Short Timeout**).
	• **Activity**—Actor or partner's port activity. **Passive** indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is **Active**. **Active** indicates the port's preference to participate in the protocol regardless of the partner's control value.

Table 23: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  ◦ **Current**—The state machine receives an LACP PDU and enters the **Current** state.
	  ◦ **Defaulted**—If no LACP PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state.
	  ◦ **Expired**—If no LACP PDU is received before the timer for the **Current** state expires once, the state machine enters the **Expired** state.
	  ◦ **Initialize**—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the **Initialize** state.
	  ◦ **LACP Disabled**—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to **LACP Disabled**. This state is similar to the **Defaulted** state, except that the port is forced to operate as an individual port.
	  ◦ **Port Disabled**—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the **Port Disabled** state.
	• **Transmit State**—Transmit state of state machine. One of the following values:
	  ◦ **Fast Periodic**—Periodic transmissions are enabled at a fast transmission rate.
	  ◦ **No Periodic**—Periodic transmissions are disabled.
	  ◦ **Periodic Timer**—Transitory state entered when the periodic timer expires.
	  ◦ **Slow Periodic**—Periodic transmissions are enabled at a slow transmission rate.
	• **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values:
	  ◦ **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.
	  ◦ **Collecting**—**Yes** indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. **No** indicates the receive function of this link is not enabled.
	  ◦ **Collecting Distributing**—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.
	  ◦ **Detached**—Process of detaching the port from the aggregator is in progress.
	  ◦ **Distributing**—**Yes** indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. **No** indicates the transmit function of this link is not enabled.
	  ◦ **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.

LACP Statistics	LACP statistics are returned when the **extensive** option is used and provides the following information: • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received.

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
 Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout    Activity

      ge-1/0/1       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/1     Partner    No   Yes    No   No   No   Yes     Fast     Passive

      ge-1/0/2       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/2     Partner    No   Yes    No   No   No   Yes     Fast     Passive

    LACP protocol:        Receive State   Transmit State      Mux State
      ge-1/0/1                CURRENT      Fast periodic      Collecting
distributing
      ge-1/0/2                CURRENT      Fast periodic      Collecting
distributing
      ge-1/0/1 (active)       CURRENT      Fast periodic      Collecting
distributing
      ge-1/0/2 (standby)      CURRENT      Fast periodic      WAITING
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx    Illegal Rx
      ge-1/0/1                   0           0            0    0
      ge-1/0/2                   0           0            0    0
```

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
    LACP State:        Role    Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/3/0        Actor     No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/3/0      Partner     No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP Protocol:   Receive State    Transmit State          Mux State
      ge-0/3/0             Current     Fast periodic  Collecting distributing
```

show interfaces mac-database (Gigabit Ethernet)

Syntax

show interfaces mac-database (ge-*fpc/pic/port* | ge-*fpc/pic/port.n*)
<mac-address *mac-address*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Description

(M Series, T Series, MX Series routers, and PTX Series Packet Transport Routers only) Display media access control (MAC) address information for the specified Gigabit Ethernet interface.

Options

ge-*fpc/pic/port*—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

ge-*fpc/pic/port.n*—Display MAC addresses that have been learned on a particular logical interface.

mac-address *mac-address*—(Optional) Display detailed MAC address statistics, including policer information.

Additional Information

On IQ2 PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

Required Privilege Level

view

List of Sample Output

show interfaces mac-database (All MAC Addresses on a Port) on page 376
show interfaces mac-database (All MAC Addresses on a Service) on page 377
show interfaces mac-database mac-address on page 377

Output Fields

Table 24 on page 374 lists the output fields for the **show interfaces mac-database** command. Output fields are listed in the approximate order in which they appear.

Table 24: show interfaces mac-database Output Fields

Field Name	Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under *Common Output Fields Description*.
Interface index	Physical interface index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Description	Description and name of the interface.

Table 24: show interfaces mac-database Output Fields *(continued)*

Field Name	Field Description
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback: **local** or **remote**.
Source filtering	Whether source filtering is configured.
Flow control	Whether flow control is enabled or disabled.
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.
Interface flags	Information about the interface. Possible values are described in the “Links Flags” section under *Common Output Fields Description*.
Link flags	Information about the link. Possible values are described in the “Device Flags” section under *Common Output Fields Description*.
Logical Interface	
Logical interface	Name of the logical interface.
Index	Logical interface index number, which reflects its initialization sequence.
SNMP ifIndex	Logical interface SNMP interface index number.
Flags	Information about the logical interface (possible values are described in the “Logical Interface Flags” section under *Common Output Fields Description*.
Encapsulation	Encapsulation on the logical interface.
MAC address, Input frames, Input bytes, Output frames, Output bytes	MAC address and corresponding number of input frames, input bytes, output frames, and output bytes.
Number of MAC addresses	Number of MAC addresses configured.

Table 24: show interfaces mac-database Output Fields *(continued)*

Field Name	Field Description

Policer Statistics	(Displayed for **mac-address** option only) Display information about policers applied to a logical interface-MAC pair.
	
	• **Policer type**—Type of policer that is out of spec with respect to the configuration. It can be one or more of the following:
	  • **Input premium**—Number of high-priority rating out-of-spec frames or bytes received.
	  • **Output premium**—Number of high-priority rating out-of-spec frames or bytes sent.
	  • **Input aggregate**—Total number of out-of-spec frames or bytes received.
	  • **Output aggregate**—Total number of out-of-spec frames or bytes sent.
	• **Discarded Frames**—Number of discarded frames.
	• **Discarded Bytes**—Number of discarded bytes.

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address         Input frames     Input bytes  Output frames    Output bytes
  00:00:00:00:00:00              1              56              0               0
  00:00:c0:01:01:02        7023810       323095260              0               0
  00:00:c0:01:01:03        7023810       323095260              0               0
  00:00:c0:01:01:04        7023810       323095260              0               0
  00:00:c0:01:01:05        7023810       323095260              0               0
  00:00:c0:01:01:06        7023810       323095260              0               0
  00:00:c0:01:01:07        7023810       323095260              0               0
  00:00:c0:01:01:08        7023809       323095214              0               0
  00:00:c0:01:01:09        7023809       323095214              0               0
  00:00:c0:01:01:0a        7023809       323095214              0               0
  00:00:c0:01:01:0b        7023809       323095214              0               0
  00:00:c8:01:01:02       30424784      1399540064       37448598      1722635508
  00:00:c8:01:01:03       30424784      1399540064       37448598      1722635508
  00:00:c8:01:01:04       30424716      1399536936       37448523      1722632058
  00:00:c8:01:01:05       30424789      1399540294       37448598      1722635508
  00:00:c8:01:01:06       30424788      1399540248       37448597      1722635462
  00:00:c8:01:01:07       30424783      1399540018       37448597      1722635462
  00:00:c8:01:01:08       30424783      1399540018       37448596      1722635416
  00:00:c8:01:01:09        8836796       406492616        8836795       406492570
  00:00:c8:01:01:0a       30424712      1399536752       37448521      1722631966
  00:00:c8:01:01:0b       30424715      1399536890       37448523      1722632058
  Number of MAC addresses : 21
```


show interfaces mac-database (All MAC Addresses on a Service)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address        Input frames     Input bytes  Output frames   Output bytes
  00:00:00:00:00:00             1              56              0              0
  00:00:c0:01:01:02       7023810       323095260              0              0
  00:00:c0:01:01:03       7023810       323095260              0              0
  00:00:c0:01:01:04       7023810       323095260              0              0
  00:00:c0:01:01:05       7023810       323095260              0              0
  00:00:c0:01:01:06       7023810       323095260              0              0
  00:00:c0:01:01:07       7023810       323095260              0              0
  00:00:c0:01:01:08       7023809       323095214              0              0
  00:00:c0:01:01:09       7023809       323095214              0              0
  00:00:c0:01:01:0a       7023809       323095214              0              0
  00:00:c0:01:01:0b       7023809       323095214              0              0
  00:00:c8:01:01:02      31016568      1426762128       38040381     1749857526
  00:00:c8:01:01:03      31016568      1426762128       38040382     1749857572
  00:00:c8:01:01:04      31016499      1426758954       38040306     1749854076
  00:00:c8:01:01:05      31016573      1426762358       38040381     1749857526
  00:00:c8:01:01:06      31016573      1426762358       38040381     1749857526
  00:00:c8:01:01:07      31016567      1426762082       38040380     1749857480
  00:00:c8:01:01:08      31016567      1426762082       38040379     1749857434
  00:00:c8:01:01:09       9428580       433714680        9428580      433714680
  00:00:c8:01:01:0a      31016496      1426758816       38040304     1749853984
  00:00:c8:01:01:0b      31016498      1426758908       38040307     1749854122
```

show interfaces mac-database mac-address

```
user@host> show interfaces mac-database xe-0/3/3 mac-address  00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
    Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes      : 202324652
    Output bytes     : 202324560
    Input frames     : 4398362
    Output frames    : 4398360
  Policer statistics:
  Policer type      Discarded frames    Discarded bytes
  Output aggregate           3992386          183649756
```

show interfaces mc-ae

Syntax

show interfaces mc-ae
<revertive-info>
<id *identifier* unit *number*>

Release Information

Command introduced in Junos OS Release 9.6.
revertive-info statetment introduced in Junos OS Release 13.3

Description

On MX Series routers with multichassis aggregated Ethernet (**ae*X***) interfaces, displaya information about the **ae*X*** interfaces.

Options

revertive-info—(Optional) Display revertive mode information for multichassis aggregated Ethernet interface.

identifier—(Optional) Identifier of the multichassis aggregated Ethernet interface.

number—(Optional) Specify the logical interface by unit number.

Required Privilege Level

view

Related Documentation

- Configuring Multichassis Link Aggregation on page 42

List of Sample Output

show interfaces mc-ae on page 379
show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers) on page 379
show interfaces mc-ae revertive-info on page 380

Output Fields

Table 25 on page 378 lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 25: show interfaces mc-ae Output Fields

Output Field Name	Field Description
Member Link	Identifiers of the configured multichassis link aggregate interfaces configured interfaces.
Local Status	Status of the local link: **active** or **standby**.
Peer Status	Status of the peer link: **active** or **standby**.
Local State	Up or down state of the local device.

Table 25: show interfaces mc-ae Output Fields *(continued)*

Output Field Name	Field Description
Peer State	Status of the local and peer links in an **active/active** bridge or VRRP over integrated routing and bridging (IRB) configuration on MX Series routers, including: Logical Interface—Aggregated Ethernet (AE) aggregate number and unit number. Topology Type—The bridge or VRRP topology type configured on the AE. Local State—Up or down state of the local device. Peer State—Up or down state of the peer device. Peer Ip/ICL-PL/State—Address, interface and state of the peer device.
Logical Interface	Identifier and unit of the multichassis aggregated Ethernet interface.
Core Facing Interface	Label: **pseudowire interface** or **Ethernet interface**.
ICL-PL	Label: **pseudowire interface** or **Ethernet interface**.
switchover mode	The configured switchover mode for the multichassis aggregated Ethernet interface: **revertive** or **non-revertive**.
switchover status	Status of the switchover if the **revert-time** statement is configured at the **[edit interfaces ae*x* mc-ae]** hierarchy level.
revert time	Revert time configured for the multichassis aggregated Ethernet interface.
switchover time remaining	Seconds left to trigger the switchover if the switchover is in progress.

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links    : ae0
  Local Status    : active
  Peer Status     : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series Routers)

```
user@host# show interfaces mc-ae ge-0/0/0.0
 Member Link                  : ae0
 Current State Machine's State: active
```

```
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
 Peer State                   : up
     Logical Interface        : ae0.0
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/ICL-PL/State     : 192.168.100.10 ge-0/0/0.0 up
```

show interfaces mc-ae revertive-info

```
user@host> show interfaces mc-ae revertive-info id 2
Member Link           : ae1
Current State Machine's State: mcae active state
Local Status                 : active
Local State                  : up
Peer Status                  : standby
Peer State                   : up
Switchover Mode              : Non Revertive
Switchover Status            : N/A
Revert Time                  : 1 Minutes
Switchover Remaining Time    : N/A
     Logical Interface        : ae1.1024
     Topology Type            : bridge
     Local State              : up
     Peer State               : up
     Peer Ip/MCP/State        : N/A
```


show 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 389
- show oam ethernet connectivity-fault-management mep-database on page 402
- show oam ethernet connectivity-fault-management mep-statistics on page 413

List of Sample Output

show oam ethernet connectivity-fault-management delay-statistics on page 383
show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 383

Output Fields

Table 26 on page 382 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 26: 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 386
show oam ethernet connectivity-fault-management forwarding-state interface on page 386
show oam ethernet connectivity-fault-management forwarding-state interface detail on page 387
show oam ethernet connectivity-fault-management forwarding-state interfaceinterface-name on page 388

Output Fields

Table 27 on page 385 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 27: 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 27: 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-3/0/0/0.0**

```
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 381
- show oam ethernet connectivity-fault-management mep-database on page 402
- show oam ethernet connectivity-fault-management mep-statistics on page 413

List of Sample Output

show oam ethernet connectivity-fault-management interfaces (trunk ports) on page 398

Output Fields Table 28 on page 390 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 28: 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 28: 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 28: 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 28: 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 28: 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@host**show oam ethernet connectivity-fault-management interfaces detail**

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



```
    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 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 400
show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 400

Output Fields Table 29 on page 399 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 29: 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 29: 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 381
- show oam ethernet connectivity-fault-management interfaces on page 389
- show oam ethernet connectivity-fault-management mep-statistics on page 413

List of Sample Output

show oam ethernet connectivity-fault-management mep-database on page 407
show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM) on page 408
show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 409

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 409
show oam ethernet connectivity-fault-management mep-database (Connection Protection TLV Configured) on page 409
show oam ethernet connectivity-fault-management mep-database on page 410
show oam ethernet connectivity-fault-management mep-database (enhanced continuity measurement) on page 411

Output Fields Table 30 on page 403 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 30: 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 30: 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 30: 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 30: 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 30: 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 381
- show oam ethernet connectivity-fault-management interfaces on page 389
- show oam ethernet connectivity-fault-management mep-database on page 402

List of Sample Output

show oam ethernet connectivity-fault-management mep-statistics (CIR counters only) on page 415
show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled) on page 417
show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR counters only) on page 418

show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled) on page 419
show oam ethernet connectivity-fault-management mep-statistics on page 421
show oam ethernet connectivity-fault-management mep-statistics remote-mep on page 422

Output Fields Table 31 on page 414 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 31: 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 31: 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 425

Output Fields

Table 32 on page 424 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 32: 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 32: 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 33 on page 426 lists the output fields for the **show oam ethernet evc** command. Output fields are listed in the approximate order in which they appear.

Table 33: 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 431
show oam ethernet link-fault-management detail on page 431

Output Fields

Table 34 on page 427 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 34: 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 34: 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 34: 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 34: 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 34: 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 35 on page 433 lists the output fields for the **show oam ethernet lmi** command. Output fields are listed in the approximate order in which they appear.

Table 35: 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 435

Output Fields Table 36 on page 435 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 36: show oam ethernet lmi statistics Output Fields

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

Sample Output

show oam ethernet lmi statistics

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

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


show protection-group ethernet-ring aps

Syntax **show protection-group ethernet-ring aps**

Release Information Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.

Description Display the status of the Automatic Protection Switching (APS) and Ring APS (RAPS) messages on an Ethernet ring.

Options This command has no options.

Required Privilege Level view

Related Documentation

- show protection-group ethernet-ring data-channel on page 439
- show protection-group ethernet-ring interface on page 441
- show protection-group ethernet-ring node-state on page 444
- show protection-group ethernet-ring statistics on page 447
- show protection-group ethernet-ring vlan on page 450

List of Sample Output

show protection-group ethernet-ring aps (EX Switches) on page 438
show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers) on page 438
show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers) on page 438
show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers) on page 438
show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers) on page 438

Output Fields Table 37 on page 437 lists the output fields for the **show protection-group ethernet-ring aps** command. Output fields are listed in the approximate order in which they appear.

Table 37: show protection-group ethernet-ring aps Output Fields

Field Name	Field Description
Ethernet Ring Name	Name configured for the Ethernet ring.
Request/State	Status of the Ethernet ring RAPS messages. • **NR**—Indicates there is no request for APS on the ring. • **SF**—Indicates there is a signal failure on the ring.
No Flush	State of the ring flushing: **No** (normal) or **Yes** (failure).
Ring Protection Link Blocked	Blocking on the ring protection link: **Yes** or **No**.

Table 37: show protection-group ethernet-ring aps Output Fields *(continued)*

Field Name	Field Description
Originator	Whether this node is the ring originator: **Yes** or **No**.
Remote Node ID	Identifier (in MAC address format) of the remote node.

Sample Output

show protection-group ethernet-ring aps (EX Switches)

```
user@switch>> show protection-group ethernet-ring aps
Ring Name     Request/state  No Flush  RPL Blocked  Originator  Remote Node ID
erp1          NR             no        yes          no          00:1F:12:30:B8:81
```

Sample Output

show protection-group ethernet-ring aps (Owner Node, Normal Operation on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               NR             No        Yes

Originator  Remote Node ID
Yes
```

show protection-group ethernet-ring aps (Ring Node, Normal Operation on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               NR             No        Yes

Originator  Remote Node ID
No          00:01:01:00:00:01
```

show protection-group ethernet-ring aps (Owner Node, Failure Condition on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               SF             No        No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

show protection-group ethernet-ring aps (Ring Node, Failure Condition on MX Routers)

```
user@host> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               SF             No        Yes

Originator  Remote Node ID
Yes         00:00:00:00:00:00
```


show protection-group ethernet-ring data-channel

Syntax

show protection-group ethernet-ring data-channel
<brief | detail>
<group-name *group-name*>

Release Information

Command introduced in Junos OS Release 10.2.

Description

On MX Series routers, display data channel information for all Ethernet ring protection groups or for a specific Ethernet ring protection group.

Options

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

group-name—(Optional) Protection group for which to display statistics. If you omit this optional field, all protection group statistics for configured groups will be displayed.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring aps on page 437
- show protection-group ethernet-ring interface on page 441
- show protection-group ethernet-ring node-state on page 444
- show protection-group ethernet-ring statistics on page 447
- show protection-group ethernet-ring vlan on page 450

List of Sample Output

show protection-group ethernet-ring data-channel on page 440

Output Fields

Table 38 on page 439 lists the output fields for the **show protection-group ethernet-ring data-channel** command. Output fields are listed in the approximate order in which they appear.

Table 38: show protection-group ethernet-ring data-channel Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet ring.
STP index	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.
Forward State	Forwarding state on the Ethernet ring. • **fowarding**—Indicates packets are being forwarded. • **discarding**—Indicates packets are being discarded.

Sample Output

show protection-group ethernet-ring data-channel

```
user@host> show protection-group ethernet-ring data-channel
Ethernet ring data channel information for protection group pg301
Interface     STP index        Forward State
ge-1/0/3      71               forwarding
ge-1/0/4      82               forwarding

Ethernet ring data channel information for protection group pg302
Interface     STP index        Forward State
ge-1/0/3      52               forwarding
ge-1/0/4      91               forwarding
```


show protection-group ethernet-ring interface

Syntax

show protection-group ethernet-ring interface

Release Information

Command introduced in Junos OS Release 9.4.

Description

Displays the status of the Automatic Protection Switching (APS) interfaces on an Ethernet ring.

Options

This command has no options.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring data-channel on page 439
- show protection-group ethernet-ring aps on page 437
- show protection-group ethernet-ring node-state on page 444
- show protection-group ethernet-ring statistics on page 447
- show protection-group ethernet-ring vlan on page 450

List of Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node) on page 442
show protection-group ethernet-ring interface (Owner Node MX Series Router) on page 442
show protection-group ethernet-ring interface (EX Series Switch Ring Node) on page 442
show protection-group ethernet-ring interface (MX Series Router Ring Node) on page 442

Output Fields

Table 39 on page 441 lists the output fields for both the EX Series switch and the MX Series router **show protection-group ethernet-ring interface** commands. Output fields are listed in the approximate order in which they appear.

Table 39: MX Series Routers show protection-group ethernet-ring interface Output Fields

Field Name	Field Description
Ethernet ring port parameters for protection group *group-name*	Output is organized by configured protection group.
Interface	Physical interfaces configured for the Ethernet ring.
Control Channel	(MX Series router only) Logical unit configured on the physical interface. • **NR**—Indicates there is no request for APS on the ring. • **SF**—Indicates there is a signal failure on the ring.
Forward State	State of the ring forwarding on the interface: **discarding** or **forwarding**.

Table 39: MX Series Routers show protection-group ethernet-ring interface Output Fields *(continued)*

Field Name	Field Description
Ring Protection Link End	Whether this interface is the end of the ring: **Yes** or **No**.
Signal Failure	Whether there a signal failure exists on the link: **Clear** or **Set**.
Admin State	State of the interface: For EX switches, **ready**, **ifl ready**, or **waiting**. For MX routers, **IFF ready** or **IFF disabled**.

Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface      Forward State  RPL End  Signal Failure  Admin State

ge-0/0/3.0     discarding     Yes        Clear           ready
ge-0/0/9.0     forwarding     No         Clear           ready
```

show protection-group ethernet-ring interface (Owner Node MX Series Router)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface     Control Channel  Forward State  Ring Protection Link End
ge-1/0/1      ge-1/0/1.1       discarding     Yes
ge-1/2/4      ge-1/2/4.1       forwarding     No

Signal Failure   Admin State
Clear          IFF ready
Clear          IFF ready
```

show protection-group ethernet-ring interface (EX Series Switch Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Ethernet ring port parameters for protection group pg101

Interface      Forward State  RPL End  Signal Failure  Admin State

ge-0/0/3.0     discarding     Yes        Clear           ready
ge-0/0/9.0     forwarding     No         Clear           ready
```

show protection-group ethernet-ring interface (MX Series Router Ring Node)

```
user@host> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface     Control Channel  Forward State  Ring Protection Link End
ge-1/2/1      ge-1/2/1.1       forwarding     No
ge-1/0/2      ge-1/0/2.1       forwarding     No
```



```
Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```

show protection-group ethernet-ring node-state

Syntax

show protection-group ethernet-ring node-state

Release Information

Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.

Description

Display the status of the Automatic Protection Switching (APS) nodes on an Ethernet ring.

Options

This command has no options.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring data-channel on page 439
- show protection-group ethernet-ring aps on page 437
- show protection-group ethernet-ring interface on page 441
- show protection-group ethernet-ring statistics on page 447
- show protection-group ethernet-ring vlan on page 450

List of Sample Output

show protection-group ethernet-ring node-state (EX Series Switch) on page 445
show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation) on page 445
show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation) on page 445
show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition) on page 446
show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition) on page 446
show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion) on page 446

Output Fields

Table 40 on page 444 lists the output fields for the **show protection-group ethernet-ring node-state** command. Output fields are listed in the approximate order in which they appear.

Table 40: show protection-group ethernet-ring node-state Output Fields

Field Name	Field Description
Ring Name/Ethernet Ring	Name configured for the Ethernet ring.

Table 40: show protection-group ethernet-ring node-state Output Fields *(continued)*

Field Name	Field Description
APS State	State of the Ethernet ring APS. • **idle**—Indicates that the ring is working in normal condition and no protection-switching request active or pending in the ring. When the ring is in idle state, it is blocked at RPL link. • **protected**—Indicates that there is a protection switch on the ring due to signal failure condition on the ring link.
Event	Events on the ring. • **NR-RB**—Indicates there is no APS request and the ring link is blocked on the ring owner node. • **NR**—Indicates there is no APS request pending in the ring. • **local SF**—Indicates there is signal failure on one or both the ring links of the node. • **remote SF**—Indicates there is signal failure on ring links of any other node of the ring. • **WTR running**—Indicates wait to restore timer is running on RPL-owner.
RPL Owner / Ring Protection Link Owner	Whether this node is the ring owner: **Yes** or **No**.
WTR Timer / Restore Timer	Restoration timer: **running** or **disabled**.
Guard Timer	Guard timer: **running** or **disabled**.
Op state / Operational State	State of the node: **Operational** or **any internal wait state**.

Sample Output

show protection-group ethernet-ring node-state (EX Series Switch)

```
user@switch> show protection-group ethernet-ring node-state
Ring Name APS State Event RPL Owner WTR Timer Guard Timer Op State
erp1      idle       NR-RB    yes     disabled    disabled    operational
```

show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         Ring Protection Link Owner
pg101             idle          NR-RB         Yes

Restore Timer  Guard Timer  Operation state
disabled       disabled      operational
```

show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
```

```
Ethernet ring     APS State     Event         Ring Protection Link Owner
pg102             idle          NR-RB         No

Restore Timer  Guard Timer  Operation state
disabled       disabled     operational
```

show protection-group ethernet-ring node-state (MX Series Router - Owner Node, Remote Signal Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         Ring Protection Link Owner
pg101             protected     remote SF           Yes

Restore Timer  Guard Timer  Operation state
disabled       disabled     operational
```

show protection-group ethernet-ring node-state (MX Series Router - Ring Node, Local Signal Failure Condition)

```
user@host> show protection-group ethernet-ring node-state
Ethernet ring     APS State     Event         Ring Protection Link Owner
pg102             protected     local SF       No

Restore Timer  Guard Timer  Operation state
disabled       disabled     operational
```

show protection-group ethernet-ring node-state detail (MX Series Router - Node state at RPL-owner after signal failure condition is cleared in the ring and before reversion)

```
user@host> show protection-group ethernet-ring node-state detail
Ethernet-Ring name              : pg_major
APS State                       : protected
Event                           : WTR running
Ring Protection Link Owner      : Yes
Restore Timer                   : running (time to expire: 269 sec)
Guard Timer                     : disabled
Operation state                 : operational
```


show protection-group ethernet-ring statistics

Syntax

show protection-group ethernet-ring statistics <group-name *group-name*>

Release Information

Command introduced in Junos OS Release 9.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.

Description

Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.

Options

group-name—Protection group for which to display statistics. In you omit this optional field, all protection group statistics for configured groups will be displayed.

Required Privilege Level

view

Related Documentation

- show protection-group ethernet-ring data-channel on page 439
- show protection-group ethernet-ring aps on page 437
- show protection-group ethernet-ring node-state on page 444
- show protection-group ethernet-ring interface on page 441
- show protection-group ethernet-ring vlan on page 450

List of Sample Output

show protection-group ethernet-ring statistics (EX Switch) on page 448
show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router) on page 448
show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router) on page 448
show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router) on page 448
show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on MX Router) on page 449

Output Fields

Table 41 on page 447 lists the output fields for the **show protection-group ethernet-ring statistics** command. Output fields are listed in the approximate order in which they appear.

Table 41: show protection-group ethernet-ring statistics Output Fields

Field Name	Field Description
Ethernet Ring Statistics for PG	Name of the protection group for which statistics are displayed.
RAPS sent	Number of Ring Automatic Protection Switching (RAPS) messages sent. (On MX Series switches only)
RAPS received	Number of RAPS messages received. (On MX Series switches only)

Table 41: show protection-group ethernet-ring statistics Output Fields *(continued)*

Field Name	Field Description
Local SF	Number of times a signal failure (SF) has occurred locally.
Remote SF	Number of times a signal failure (SF) has occurred anywhere else on the ring.
NR event	Number of times a No Request (NR) event has occurred on the ring.
NR-RB event	Number of times a No Request, Ring Blocked (NR-RB) event has occurred on the ring.

Sample Output

show protection-group ethernet-ring statistics (EX Switch)

```
user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1       2       1       2        3
```

show protection-group ethernet-ring statistics (Owner Node, Normal Operation on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                         : 1
RAPS received                     : 0
Local SF happened                 : 0
Remote SF happened                : 0
NR event happened                 : 0
NR-RB event happened              : 1
```

show protection-group ethernet-ring statistics (Ring Node, Normal Operation on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                         : 0
RAPS received                     : 1
Local SF happened                 : 0
Remote SF happened                : 0
NR event happened                 : 0
NR-RB event happened              : 1
```

show protection-group ethernet-ring statistics (Owner Node, Failure Condition on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                         : 1
RAPS received                     : 1
Local SF happened                 : 0
Remote SF happened                : 1
NR event happened                 : 0
NR-RB event happened              : 1
```


show protection-group ethernet-ring statisitics (Ring Node, Failure Condition on MX Router)

```
user@host> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg102
RAPS sent                         : 1
RAPS received                     : 1
Local SF happened                 : 1
Remote SF happened                : 0
NR event happened                 : 0
NR-RB event happened              : 1
```

show protection-group ethernet-ring vlan

Syntax

show protection-group ethernet-ring vlan
<brief | detail>
<group-name *group-name*>

Release Information Command introduced in Junos OS Release 10.2.

Description On MX Series routers, display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

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

group-name—(Optional) Protection group for which to display details such as data channel interfaces, vlan, and bridge-domain. If you omit this optional field, details for all configured protection groups will be displayed.

Required Privilege Level view

Related Documentation

- show protection-group ethernet-ring aps on page 437
- show protection-group ethernet-ring data-channel on page 439
- show protection-group ethernet-ring interface on page 441
- show protection-group ethernet-ring node-state on page 444
- show protection-group ethernet-ring statistics on page 447

List of Sample Output

show protection-group ethernet-ring vlan on page 451
show protection-group ethernet-ring vlan brief on page 451
show protection-group ethernet-ring vlan detail on page 451
show protection-group ethernet-ring vlan group-name vkm01 on page 451

Output Fields Table 42 on page 450 lists the output fields for the **show protection-group ethernet-ring vlan** command. Output fields are listed in the approximate order in which they appear.

Table 42: show protection-group ethernet-ring vlan Output Fields

Field Name	Field Description
Interface	Name of the interface configured for the Ethernet protection ring.
Vlan	Name of the VLAN associated with the interface configured for the Ethernet protection ring.
STP index	The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior.

Table 42: show protection-group ethernet-ring vlan Output Fields *(continued)*

Field Name	Field Description
Bridge Domain	Name of the bridge domain that is associated with the VLAN configured for the Ethernet protection ring.

Sample Output

show protection-group ethernet-ring vlan

```
user@host> show protection-group ethernet-ring vlan
Ethernet ring IFBD parameters for protection group vkm01

Interface     Vlan     STP Index    Bridge Domain
ge-2/0/8      100      130          default-switch/bd100
ge-2/0/4      100      126          default-switch/bd100
```

show protection-group ethernet-ring vlan brief

```
user@host> show protection-group ethernet-ring vlan brief
Ethernet ring IFBD parameters for protection group vkm01

Interface     Vlan     STP Index    Bridge Domain
ge-2/0/8      100      130          default-switch/bd100
ge-2/0/4      100      126          default-switch/bd100
```

show protection-group ethernet-ring vlan detail

```
user@host> show protection-group ethernet-ring vlan detail
Ethernet ring IFBD parameters for protection group vkm01

Interface name                 : ge-2/0/8
Vlan                           : 100
STP index                      : 130
Bridge Domain                  : default-switch/bd100

Interface name                 : ge-2/0/4
Vlan                           : 100
STP index                      : 126
Bridge Domain                  : default-switch/bd100
```

show protection-group ethernet-ring vlan group-name vkm01

```
user@host> show protection-group ethernet-ring vlan vkm01
Ethernet ring IFBD parameters for protection group vkm01

Interface     Vlan     STP Index    Bridge Domain
ge-2/0/8      100      130          default-switch/bd100
ge-2/0/4      100      126          default-switch/bd100
```

PART 4

Troubleshooting

CHAPTER 6

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 457

Output Fields

Table 43 on page 456 lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 43: 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 43: 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
```


Tracing PPPoE Operations

The Junos OS trace feature tracks PPPoE operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file called **pppoed** located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located.
2. When the file **pppoed** reaches 128 kilobytes (KB), it is renamed **pppoed.0**, then **pppoed.1**, and finally **pppoed.2**, until there are three trace files. Then the oldest trace file (**pppoed.2**) is overwritten.

   You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPPoE tracing operations:

1. Specify that you want to configure tracing options.

   ```
   [edit protocols pppoe]
   user@host# edit traceoptions
   ```
2. (Optional) Configure the name for the file used for the trace output.
3. (Optional) Configure the number and size of the log files.
4. (Optional) Configure access to the log file.
5. (Optional) Configure a regular expression to filter logging events.
6. (Optional) Configure flags to filter the operations to be logged.

Optional PPPoE traceoptions operations are described in the following sections:

Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is **pppoed**. You can specify a different name with the **file** option.

Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

- ```
  [edit protocols pppoe traceoptions]
  user@host# set flag authentication
  ```

Related Documentation

- *PPPoE Overview*
- *Ethernet Interfaces*

Troubleshooting PPPoE Service Name Tables

Problem A misconfiguration of a PPPoE service name table can prevent PPPoE services from being properly activated. Configuration options for PPPoE service name tables are simple, which should simplify discovering where a misconfiguration exists. PPPoE clients cannot


  ```

connect if the service name table contains no match for the service name tag carried in the PADI packet.

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

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

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

```
 Active PPPoE sessions: 2
    PacketType   Sent   Received
      PADI   0   16
      PADO   16   0
      PADR   0   16
      PADS   16   0
      PADT   0   0
      Service name error   0   0
      AC system error   0   0
      Generic error   0   0
      Malformed packets   0   0
      Unknown packets   0   0
```

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

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

```
Active PPPoE sessions: 0
    PacketType                      Sent         Received
      PADI                             0              100
      PADO                             0                0
      PADR                             0                0
      PADS                             0                0
      PADT                             0                0
      Service name error               0                0
      AC system error                  0                0
      Generic error                    0                0
      Malformed packets                0                0
      Unknown packets                  0                0
```

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

`user@host>` **monitor traffic interface ge-2/0/3.1 print-hex print-ascii**

```
 Listening on ge-2/0/3.1, capture size 96 bytes

 11:49:41.436682  In PPPoE PADI [Service-Name "Service1"] [Host-Uniq UTF8]
[TAG-0x120 UTF8] [Vendor-Specific UTF8]
```



```
0x0000   ffff ffff ffff 0090 1a42 0ac1 8100 029a     .........B......
0x0010   8863 1109 0000 00c9 0101 0008 5365 7276     .c..........Serv
0x0020   6963 6531 0103 0004 1200 9c43 0120 0002     ice1.......C....
0x0030   044a 0105 00ab 0000 0de9 0124 783a 3132     .J.........$x:12
0x0040   3030 3963                                   009c
```

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

Cause Typical misconfigurations appear in the service name table configurations.

Solution Use the appropriate statements to correct the misconfiguration.

Related Documentation

- *Configuring PPPoE Service Name Tables*
- *show pppoe service-name-tables*
- *show pppoe statistics*
- *show pppoe underlying-interfaces*
- *PPPoE Overview*
- *Ethernet Interfaces*

Verifying a PPPoE Configuration

Purpose You can use show commands to display and verify the PPPoE configuration.

Action To verify a PPPoE configuration, you can issue the following operational mode commands:

- **show interfaces at-*fpc/pic/port* extensive**
- **show interfaces pp0**
- **show pppoe interfaces**
- **show pppoe version**
- **show pppoe service-name-tables**
- **show pppoe sessions**
- **show pppoe statistics**
- **show pppoe underlying-interfaces**

For more information about these operational mode commands, see CLI Explorer.

Related Documentation

- *PPPoE Overview*
- *Ethernet Interfaces*

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 44 on page 464 shows the loopback modes supported on the various interface types.

Table 44: 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 44: 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 0.152 standard)
  pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.152 standard)
  pseudo-2e20-o151     Pattern is 2^20 - 1 (per 0.151 standard)
  pseudo-2e20-o153     Pattern is 2^20 - 1 (per 0.153 standard)
        ...
```

For specific hierarchy information, see the individual interface types.

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

```
pseudo-2e11-o152     Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151     Pattern is 2^23 (per 0.151 standard)
```

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

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

```
all-ones-repeating   Repeating one bits
all-zeros-repeating  Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e11-o152     Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7           Pattern is 2^7 - 1
pseudo-2e9-o153      Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4     1 bit in 4 is set
repeating-1-in-8     1 bit in 8 is set
repeating-3-in-24    3 bits in 24 are set
```

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

NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating   Repeating one bits
all-zeros-repeating  Repeating zero bits
alternating-double-ones-zeros  Alternating pairs of ones and zeros
alternating-ones-zeros  Alternating ones and zeros
pseudo-2e9-o153      Pattern is 2^9  -1 (per 0.153 (511 type) standard)
pseudo-2e11-o152    Pattern is 2^11 -1 (per 0.152 and 0.153 (2047 type)
 standards)
pseudo-2e15-o151     Pattern is 2^15 -1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 -1 (per 0.151 standard)
pseudo-2e20-o153     Pattern is 2^20 -1 (per 0.153 standard)
pseudo-2e23-o151     Pattern is 2^23 -1 (per 0.151 standard)
repeating-1-in-4     1 bit in 4 is set
repeating-1-in-8     1 bit in 8 is set
repeating-3-in-24    3 bits in 24 are set
```

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

NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

```
all-ones-repeating            Repeating one bits
all-zeros-repeating           Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros        Alternating ones and zeros
repeating-1-in-4              1 bit in 4 is set
repeating-1-in-8              1 bit in 8 is set
repeating-3-in-24             3 bits in 24 are set
pseudo-2e9-o153               Pattern is 2^9 - 1 (per 0.153 standard)
pseudo-2e11-o152              Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151              Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151              Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153              Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151              Pattern is 2^23 (per 0.151 standard)
```

Table 45 on page 468 shows the BERT capabilities for various interface types.

Table 45: 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

R

S

T

V