Junos® OS for EX Series Ethernet Switches

Ethernet Switching on EX9200 Switches

Release
12.3

Published: 2013-04-04

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

Junos® OS for EX Series Ethernet Switches Ethernet Switching on EX9200 Switches
Release 12.3

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration 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)	Enclose 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)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web 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 J-Web 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: https://www.juniper.net/alerts/

- 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

Bridging

Layer 2 VLANs Overview

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface, and virtual switches that isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space.

A VLAN is a set of logical ports that share the same flooding or broadcast characteristics and span one or more ports of multiple devices.

You can configure one or more VLANs to perform Layer 2 bridging. Thus, MX Series routers or EX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN. Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.

You can also group one or more VLANs within a single instance, or virtual switch. Multiple virtual switches, each of which operates independently of other virtual switches on the device, are supported. Virtual switches isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space. Thus, each virtual switch can participate in a different Layer 2 network.

VLANs provide support for a Layer 2 trunk port. A Layer 2 trunk interface enables you to configure a single logical interface to represent multiple VLANs on a physical interface. You can configure a set of VLANs and VLAN identifiers that are automatically associated with one or more Layer 2 trunk interfaces. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A Layer 2 trunk interface

also supports IRB within a VLAN. In addition, you can configure Layer 2 learning and forwarding properties that apply to the entire set of VLANs.

You can configure VPLS ports in a virtual switch instead of a dedicated routing instance of type **vpls** so that the logical interfaces of the Layer 2 VLANs in the virtual switch can handle VPLS routing instance traffic. Packets received on a Layer 2 trunk interface are forwarded within a VLAN that has the same VLAN identifier.

Layer 2 Learning and Forwarding Overview

You can configure Layer 2 MAC address and VLAN learning and forwarding properties in support of Layer 2 bridging. Unicast media access control (MAC) addresses are learned to avoid flooding the packets to all the ports in a VLAN. A source MAC entry is created in its source and destination MAC tables for each MAC address learned from packets received on ports that belong to the VLAN.

By default, Layer 2 address learning is enabled. You can disable MAC learning for a device or for a specific VLAN or logical interfaces. You can also configure the following Layer 2 forwarding properties:

- Timeout interval for MAC entries
- MAC accounting
- A limit to the number of MAC addresses learned from the logical interfaces

For more information about how to configure VLANs and virtual switches, see "Configuring a VLAN" on page 18 and "Configuring a Layer 2 Virtual Switch" on page 35.

Layer 2 Learning and Forwarding for VLANs Overview

When you configure a VLAN, Layer 2 address learning is enabled by default. The VLAN learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the VLAN. Each VLAN creates a source MAC entry in its source and destination MAC tables for each source MAC address learned from packets received on the ports that belong to the VLAN.

NOTE: **Traffic is not flooded back onto the interface on which it was received. However, because this "split horizon" occurs at a late stage, the packet statistics displayed by commands such as show interfaces queue will include flood traffic.**

You can optionally disable MAC learning either for the entire device or for a specific VLAN or logical interface. You can also configure the following Layer 2 learning and forwarding properties:

- Static MAC entries for logical interfaces only
- Limit to the number of MAC addresses learned from a specific logical interface or from all the logical interfaces in a VLAN

- Size of the MAC address table for the VLAN
- MAC accounting for a VLAN

Related Documentation
- Layer 2 Learning and Forwarding Overview on page 4

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

Layer 2 learning is enabled by default. A set of VLANs, configured to function as a switch with a Layer 2 trunk port, learns unicast media access control (MAC) addresses to avoid flooding packets to the trunk port.

NOTE: **Traffic is not flooded back onto the interface on which it was received. However, because this "split horizon" occurs at a late stage, the packet statistics displayed by commands such as show interfaces queue will include flood traffic.**

You can optionally disable Layer 2 learning for the entire set of VLANs as well as modify the following Layer 2 learning and forwarding properties:

- Limit the number of MAC addresses learned from the Layer 2 trunk port associated with the set of VLANs
- Modify the size of the MAC address table for the set of VLANs
- Enable MAC accounting for the set of VLANs

Related Documentation
- Layer 2 Learning and Forwarding Overview on page 4

Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

For a VLAN that is performing Layer 2 switching only, you do not have to specify a VLAN identifier.

For a VLAN that is performing Layer 3 IP routing, you must specify either a VLAN identifier or dual VLAN identifier tags.

For a VPLS routing instance, you must specify either a VLAN identifier or dual VLAN identifier tags.

Related Documentation
- Layer 2 Learning and Forwarding Overview on page 4

CHAPTER 2

MVRP

Understanding Multiple VLAN Registration Protocol (MVRP)

You can configure Multiple VLAN Registration Protocol (MVRP) on Juniper Networks MX Series routers and EX Series switches. The primary purpose of MVRP is to manage dynamic VLAN registration in switching networks. In managing dynamic VLAN registration, MVRP also prunes VLAN information.

MVRP is an Layer 2 application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

MVRP was created by IEEE as a replacement application for GVRP. MVRP and GVRP cannot be run concurrently to share VLAN information in a switching network.

This topic describes:

How MVRP Works

The VLAN registration information sent by MVRP protocol data units (PDUs) includes the current VLANs membership—that is, which routers are members of which VLANs—and which router interfaces are in which VLAN. MVRP shares all information in the PDU with all routers participating in MVRP in the switching network.

MVRP stays synchronized using these PDUs. The routers in the network participating in MVRP receive these PDUs during state changes and update their MVRP states accordingly.

MVRP timers dictate when PDUs can be sent and when routers receiving MVRP PDUs can update their MVRP information.

VLAN information is distributed as part of the MVRP message exchange process and can be used to dynamically create VLANs, which are VLANs created on one switch and propagated to other routers as part of the MVRP message exchange process. Dynamic VLAN creation using MVRP is enabled by default but can be disabled.

As part of ensuring that VLAN membership information is current, MVRP removes routers and interfaces from the VLAN information when they become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Basics of MVRP

MVRP is disabled by default. You can configure MVRP router interfaces to participate in MVRP for the switching network. MVRP can only be enabled on trunk interfaces, and dynamic VLAN configuration through MVRP is enabled by default when MVRP is enabled.

MVRP Registration Modes

The MVRP registration mode defines whether an interface does or does not participate in MVRP.

The following MVRP registration modes are configurable:

- forbidden—The interface does not register or declare VLANS (except statically configured VLANs).
- normal—The interface accepts MVRP messages and participates in MVRP. This is the default registration mode setting.
- restricted—The interface—The interface ignores all MVRP JOIN messages received for VLANs that are not statically configured on the interface.

MRP Timers

MVRP registration and updates are controlled by timers that are part of the MRP protocol. These timers are set on a per-interface basis and define when MVRP PDUs can be sent and when MVRP information can be updated on a switch.

The following timers are used to control the operation of MVRP:

- Join timer—Controls the interval for the next MVRP PDU transmit opportunity.
- Leave timer—Controls the period of time that an interface on the switch waits in the Leave state before changing to the unregistered state.

- LeaveAll timer—Controls the frequency with which the interface generates LeaveAll messages.

BEST PRACTICE: **Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.**

MRP VLAN Messages

MVRP uses MRP messages to register and declare MVRP states for a switch and to inform the switching network that a switch is leaving MVRP. These messages are communicated as part of the PDU to communicate the state of a particular switch interface on the switching network to the other switches in the network.

The following messages are communicated for MVRP:

- Empty—VLAN information is not being declared and is not registered.
- In—VLAN information is not being declared but is registered.
- JoinEmpty—VLAN information is being declared but not registered.
- JoinIn—VLAN information is being declared and is registered.
- Leave—VLAN information that was previously registered is being withdrawn.
- LeaveAll—All registrations will be de-registered. Participants that want to participate in MVRP will need to re-register.
- New—VLAN information is new and possibly not previously registered.

MVRP Limitations

The following limitations apply when configuring MVRP:

- MVRP works with Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP), but not with VLAN Spanning Tree Protocol (VSTP).
- MVRP is allowed only on single tagged trunk ports.
- MVRP is not allowed if a physical interface has more than one logical interface.
- MVRP is only allowed if a logical has one trunk interface (unit 0).

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Controlling the Management State of a VLAN in MVRP Configurations (CLI Procedure)
- Verifying That MVRP Is Working Correctly on page 129

CHAPTER 3

Proxy ARP

Restricted and Unrestricted Proxy ARP Overview

By default, the Junos OS responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet Interfaces, you can configure the router or switches to proxy-reply to the ARP requests using the restricted or unrestricted proxy ARP configuration.

You might want to configure restricted or unrestricted proxy ARP for routers that act as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.

NOTE: From Junos OS Release 10.0 onward, Junos OS does not respond to proxy ARP requests with the default route 0.0.0.0. This behavior is in compliance with RFC 1027.

Restricted Proxy ARP

Restricted proxy ARP enables the router or switch to respond to the ARP requests in which the physical networks of the source and target are not the same and the router or switch has an active route to the target address in the ARP request. The router does not reply if the target address is on the same subnet and the same interface as the ARP requestor.

Unrestricted Proxy ARP

Unrestricted proxy ARP enables the router or switch to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.

WARNING: If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2

reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.

NOTE: While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

Topology Considerations for Unrestricted Proxy ARP

In most situations, you should not configure the router or switch to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used. Figure 1 on page 12 and Figure 2 on page 13 show examples of situations in which you might want to configure unrestricted proxy ARP.

In Figure 1 on page 12, the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In Figure 2 on page 13, the Broadband Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

Figure 1: Edge Device Case for Unrestricted Proxy ARP

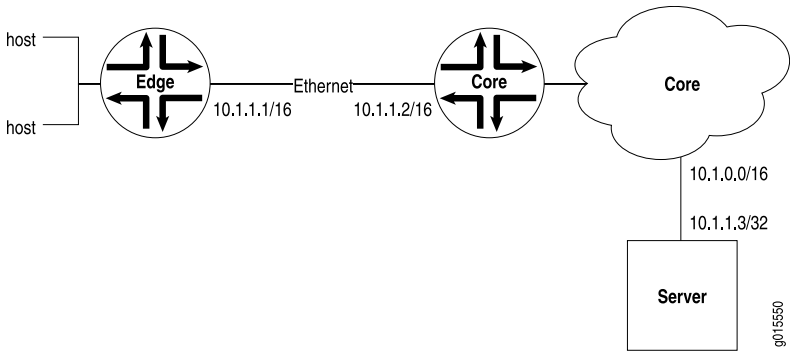

Figure 2: Core Device Case for Unrestricted Proxy ARP

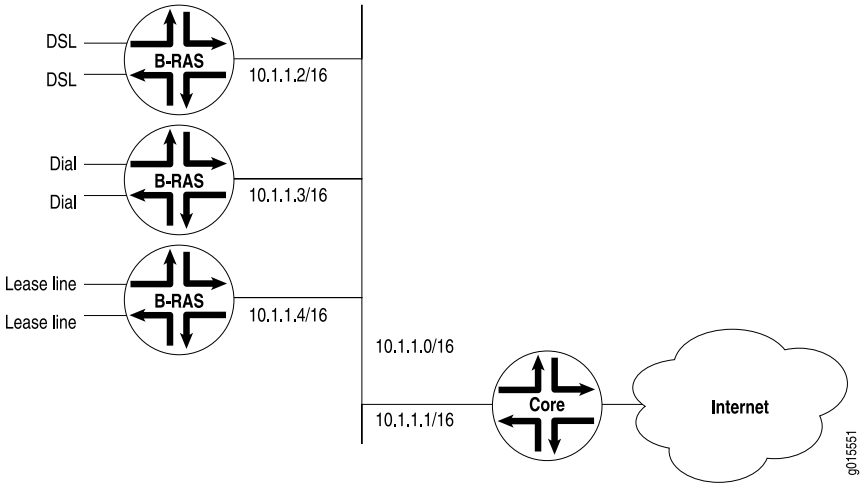

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 38
- Junos® OS Ethernet Interfaces

PART 2

Configuration

CHAPTER 4

Configuration Tasks

Configuring a VLAN

A VLAN must include a set of logical interfaces that participate in Layer 2 learning and forwarding. You can optionally configure a VLAN identifier and a Layer 3 interface for the VLAN to also support Layer 3 IP routing.

To enable a VLAN, include the following statements:

```
[edit]
vlans {
  vlan-name {
    domain-type bridge:
    interface interface-name;
    l3-interface interface-name;
    vlan-id (none | all | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer number inner number);
  }
}
```

You cannot use the slash (/) character in VLAN names. If you do, the configuration does not commit and an error is generated.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** or **all** options.

To include one or more logical interfaces in the VLAN, specify an ***interface-name*** for an Ethernet interface you configured at the **[edit interfaces]** hierarchy level.

NOTE: A maximum of 4096 active logical interfaces are supported for a VLAN or on each mesh group in a virtual private LAN service (VPLS) instance configured for Layer 2 bridging.

By default, each VLAN maintains a Layer 2 forwarding database that contains media access control (MAC) addresses learned from packets received on the ports that belong to the VLAN. You can modify Layer 2 forwarding properties, for example, disabling MAC learning for the entire system or a VLAN, adding static MAC addresses for specific logical interfaces, and limiting the number of MAC addresses learned by the entire system, the VLAN, or a logical interface.

You can also configure spanning tree protocols to prevent forwarding loops.

Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

You can configure VLAN identifiers for a VLAN or a VPLS routing instance in the following ways:

- By using either the **vlan-id** statement or the **vlan-tags** statement to configure a normalizing VLAN identifier. This topic describes how normalizing VLAN identifiers are processed and translated in a VLAN or a VPLS routing instance.

- By using the **input-vlan-map** and the **output-vlan-map** statements at the **[edit interfaces** *interface-name* **unit** *logic-unit-number***]** or **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logic-unit-number***]** hierarchy level to configure VLAN mapping.

The **vlan-id** and **vlan-tags** statements are used to specify the normalizing VLAN identifier under the VLAN or VPLS routing instance. The normalizing VLAN identifier is used to perform the following functions:

- Translate, or normalize, the VLAN tags of packets received into a learn VLAN identifier.
- Create multiple learning domains that each contain a learn VLAN identifier. A learning domain is a MAC address database to which MAC addresses are added based on the learn VLAN identifier.

NOTE: You cannot configure VLAN mapping using the **input-vlan-map** and **output-vlan-map** statements if you configure a normalizing VLAN identifier for a VLAN or VPLS routing instance using the **vlan-id** or **vlan-tags** statements.

To configure a VLAN identifier for a VLAN, include either the **vlan-id** or the **vlan-tags** statement at the **[edit interfaces** *interface-name* **unit** *logic-unit-number***]** or **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logic-unit-number***]** hierarchy level, and then include that logical interface in the VLAN configuration.

For a VPLS routing instance, include either the **vlan-id** or **vlan-tags** statement at the **[edit interfaces** *interface-name* **unit** *logic-unit-number***]** or **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logic-unit-number***]** hierarchy level, and then include that logical interface in the VPLS routing instance configuration.

NOTE: For a single VLAN or VPLS routing instance, you can include either the **vlan-id** or the **vlan-tags** statement, but not both. If you do not configure a **vlan-id**, **vlan-tags**, or **vlan-id-list [** *vlan-id-numbers* **]** for the VLAN or the VPLS routing instance, the Layer 2 packets received are forwarded to the outbound Layer 2 interface without having the VLAN tag modified unless an **output-vlan-map** is configured on the Layer 2 interface. This results in a frame being forwarded to a Layer 2 interface with a VLAN tag that is different from what is configured for the Layer 2 interface. Note that a frame received from the Layer 2 interface is still required to match the VLAN tag(s) specified in the interface configuration. The invalid configuration may cause a Layer 2 loop to occur.

The VLAN tags associated with the inbound logical interface are compared with the normalizing VLAN identifier. If the tags are different, they are rewritten as described in Table 3 on page 22. The source MAC address of a received packet is learned based on the normalizing VLAN identifier.

NOTE: You do not have to specify a VLAN identifier for a VLAN that is performing Layer 2 switching only. To support Layer 3 IP routing, you must specify either a VLAN identifier or a pair of VLAN tags. However, you cannot specify the same VLAN identifier for more than one VLAN within a routing instance. Each VLAN must have a unique VLAN identifier.

If the VLAN tags associated with the outbound logical interface and the normalizing VLAN identifier are different, the normalizing VLAN identifier is rewritten to match the VLAN tags of the outbound logical interface, as described in Table 4 on page 23.

For the packets sent over the VPLS routing instance to be tagged by the normalizing VLAN identifier, include one of the following configuration statements:

- **vlan-id** *number* to tag all packets that are sent over the VPLS virtual tunnel (VT) interfaces with the VLAN identifier.
- **vlan-tags outer** *number* **inner** *number* to tag all packets sent over the VPLS VT interfaces with dual outer and inner VLAN tags.

Use the **vlan-id none** statement to have the VLAN tags removed from packets associated with an inbound logical interface when those packets are sent over VPLS VT interfaces. Note that those packets might still be sent with other customer VLAN tags.

The **vlan-id all** statement enables you to configure bridging for several VLANs with a minimum amount of configuration. Configuring this statement creates a learning domain for:

- Each inner VLAN, or learn VLAN, identifier of a logical interface configured with two VLAN tags
- Each VLAN, or learn VLAN, identifier of a logical interface configured with one VLAN tag

We recommend that you do not use customer VLAN IDs in a VPLS routing instance because customer VLAN IDs are used for learning only.

You should use the service VLAN ID in a VPLS routing instance, as in the following configuration:

```
[edit]
interface ge-1/1/1 {
   vlan-tagging;
   unit 1 {
      vlan-id s1; /* Service vlan */
      encapsulation vlan-vpls;
      input-vlan-map pop; /* Pop the service vlan on input */
      output-vlan-map push; /* Push the service vlan on output */
   }
}
interface ge-1/1/2 {
   encapsulation ethernet-vpls;
   unit 0;
```



```
}
routing-instances {
  V1 {
    instance-type vpls;
    vlan-id all;
    interface ge-1/1/1.1;
    interface ge-1/1/2.0;
  }
}
```

NOTE: If you configure the **vlan-id all** statement in a VPLS routing instance, we recommend using the **input-vlan-map pop** and **output-vlan-map push** statements on the logical interface to pop the service VLAN ID on input and push the service VLAN ID on output and in this way, limit the impact of double-tagged frames on scaling. You cannot use the native **vlan-id** statement when the **vlan-id all** statement is included in the configuration.

The **vlan-id-list [** ***vlan-id-numbers*** **]** statement enables you to configure bridging for multiple VLANs on a trunk interface. Configuring this statement creates a learning domain for:

- Each VLAN listed: **vlan-id-list [100 200 300]**
- Each VLAN in a range: **vlan-id-list [100-200]**
- Each VLAN in a list and range combination: **vlan-id-list [50, 100-200, 300]**

The following steps outline the process for bridging a packet received over a Layer 2 logical interface when you specify a normalizing VLAN identifier using either the **vlan-id** ***number*** or **vlan-tags** statement for a VLAN or a VPLS routing instance:

1. When a packet is received on a physical port, it is accepted only if the VLAN identifier of the packet matches the VLAN identifier of one of the logical interfaces configured on that port.
2. The VLAN tags of the received packet are then compared with the normalizing VLAN identifier. If the VLAN tags of the packet are different from the normalizing VLAN identifier, the VLAN tags are rewritten as described in Table 3 on page 22.
3. If the source MAC address of the received packet is not present in the source MAC table, it is learned based on the normalizing VLAN identifier.
4. The packet is then forwarded toward one or more outbound Layer 2 logical interfaces based on the destination MAC address. A packet with a known unicast destination MAC address is forwarded only to one outbound logical interface. For each outbound Layer 2 logical interface, the normalizing VLAN identifier configured for the VLAN or VPLS routing instance is compared with the VLAN tags configured on that logical interface. If the VLAN tags associated with an outbound logical interface do not match the normalizing VLAN identifier configured for the VLAN or VPLS routing instance, the VLAN tags are rewritten as described in Table 4 on page 23.

The tables below show how VLAN tags are applied for traffic sent to and from the VLAN, depending on how the **vlan-id** and **vlan-tags** statements are configured for the VLAN and

on how identifiers are configured for the logical interfaces in a VLAN or VPLS routing instance. Depending on your configuration, the following rewrite operations are performed on VLAN tags:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack.
- **pop-pop**—Remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—Remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame.
- **swap**—Replace the VLAN tag of the frame.
- **push**—Add a new VLAN tag to the top of the VLAN stack.
- **push-push**—Push two VLAN tags in front of the frame.
- **swap-push**—Replace the VLAN tag of the frame and add a new VLAN tag to the top of the VLAN stack.
- **swap-swap**—Replace both the outer and inner VLAN tags of the frame.

Table 3 on page 22 shows specific examples of how the VLAN tags for packets sent to the VLAN are processed and translated, depending on your configuration. "–" means that the statement is not supported for the specified logical interface VLAN identifier. "No operation" means that the VLAN tags of the received packet are not translated for the specified input logical interface.

Table 3: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a VLAN

VLAN Identifier of Logical Interface	VLAN Configurations for a VLAN			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	No operation	push 200	–	push 100, push 300
200	pop 200	No operation	No operation	swap 200 to 300, push 100
1000	pop 1000	swap 1000 to 200	No operation	swap 1000 to 300, push 100
vlan-tags outer 2000 inner 300	pop 2000, pop 300	pop 2000, swap 300 to 200	pop 2000	swap 2000 to 100
vlan-tags outer 100 inner 400	pop 100, pop 400	pop 100, swap 400 to 200	pop 100	swap 400 to 300
vlan-id-range 10-100	–	–	No operation	–
vlan-tags outer 200 inner-range 10-100	–	–	pop 200	–

Table 4 on page 23 shows specific examples of how the VLAN tags for packets sent from the VLAN are processed and translated, depending on your configuration. "–" means that the statement is not supported for the specified logical interface VLAN identifier. "No operation" means that the VLAN tags of the outbound packet are not translated for the specified output logical interface.

Table 4: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a VLAN

VLAN Identifier of Logical Interface	VLAN Configurations for a VLAN			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	no operation	pop 200	–	pop 100, pop 300
200	push 200	No operation	No operation	pop 100, swap 300 to 200
1000	push 1000	swap 200 to 1000	No operation	pop 100, swap 300 to 1000
vlan-tags outer 2000 inner 300	push 2000, push 300	swap 200 to 300, push 2000	push 2000	swap 100 to 2000
vlan-tags outer 100 inner 400	push 100, push 400	swap 200 to 400, push 100	push 100	swap 300 to 400
vlan-id-range 10-100	–	–	No operation	–
vlan-tags outer 200 inner-range 10-100	–	–	push 200	–

Configuring Integrated Routing and Bridging for VLANs

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has an IRB interface configured. You configure a logical routing interface by specifying **irb** as an interface name at the **[edit interfaces]** hierarchy level and including that interface in the VLAN.

NOTE: You can include only one Layer 3 interface in a VLAN.

To configure a VLAN with IRB support, include the following statements:

```
[edit]
vlans {
    vlan-name {
        domain-type bridge;
        interface interface-name;
        l3-interface interface-name;
```

```
        vlan-id (none | number);
        vlan-tags outer number inner number;
    }
}
```

For each VLAN that you configure, specify a ***vlan-name***. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** option.

NOTE: If you configure a Layer 3 interface to support IRB in a VLAN, you cannot use the **all** option for the **vlan-id** statement.

The **vlan-tags** statement enables you to specify a pair of VLAN identifiers; an **outer** tag and an **inner** tag.

NOTE: For a single VLAN, you can include either the **vlan-id** statement or the **vlan-tags** statement, but not both.

To include one or more logical interfaces in the VLAN, specify the ***interface-name*** for each Ethernet interface to include that you configured at the **[edit interfaces]** hierarchy level.

NOTE: A maximum of 4096 active logical interfaces are supported for a VLAN or on each mesh group in a VPLS routing instance configured for Layer 2 bridging.

To associate a Layer 3 interface with a VLAN, include the **l3-interface** ***interface-name*** statement and specify an ***interface-name*** you configured at the **[edit interfaces irb]** hierarchy level. You can configure only one Layer 3 interface for each VLAN.

IRB interfaces are supported for multicast snooping.

In multihomed VPLS configurations, you can configure VPLS to keep a VPLS connection up if only an IRB interface is available by configuring the **irb** option for the **connectivity-type** statement at the **[edit routing-instances** ***routing-instance-name*** **protocols vpls]** hierarchy level. The **connectivity-type** statement has the **ce** and **irb** options. The **ce** option is the default and specifies that a CE interface is required to maintain the VPLS connection. By default, if only an IRB interface is available, the VPLS connection is brought down.

NOTE: When you configure IRB interfaces in more than one logical system on a device, all of the IRB logical interfaces share the same MAC address.

Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port

You can configure a set of VLANs that are associated with a Layer 2 trunk port. The set of VLANs function as a switch. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A trunk interface also provides support for IRB, which provides support for Layer 2 bridging and Layer 3 IP routing on the same interface.

To configure a Layer 2 trunk port and set of VLANs, include the following statements:

```
[edit interfaces]
interface-name {
  unit number {
    family ethernet-switching {
      interface-mode access;
      vlan-members (vlan-name | vlan-tag);
    }
  }
}
interface-name {
  native-vlan-id number;
  unit number {
    family ethernet-switching {
      interface-mode trunk;
      vlan-members (vlan-name | vlan-tag);
    }
  }
}
[edit vlans ]
vlan-name {
  vlan-id number;
  vlan-id-list [ vlan-id-numbers ];
  ....
}
```

You must configure a VLAN and VLAN identifier for each VLAN associated with the trunk interface. You can configure one or more trunk or access interfaces at the **[edit interfaces]** hierarchy level. An access interface enables you to accept packets with no VLAN identifier.

Disabling MAC Learning for a VLAN or Logical Interface

You can disable MAC learning for all logical interfaces in a specified VLAN, or for a specific logical interface in a VLAN. Disabling dynamic MAC learning prevents the specified interfaces from learning source MAC addresses.

To disable MAC learning for all logical interfaces in a VLAN in a virtual switch, include the **no-mac-learning** statement at the **[edit vlans *vlan-name* switch-options]** hierarchy level:

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    switch-options {
```

```
        no-mac-learning;
      }
    }
  }
```

To disable MAC learning for a specific logical interface in a VLAN, include the **no-mac-learning** statement at the **[edit vlans *vlan-name* switch-options interface *interface-name*]** hierarchy level.

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    switch-options {
      interface interface-name {
        no-mac-learning;
      }
    }
  }
}
```

NOTE: When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the VLAN.

NOTE: When you gather interfaces into a VLAN, the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* ether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit vlans *vlan-name* switch-options interface *interface-name*]** hierarchy level to disable MAC learning on an interface in a VLAN.

NOTE: When MAC learning is disabled for a VPLS routing instance, traffic is not load balanced and only one of the equal-cost next hops is used.

Disabling MAC Learning for a Set of VLANs

You can disable MAC learning for a set of VLANs. Disabling dynamic MAC learning prevents the Layer 2 trunk port associated with the set of VLANs from learning source and destination MAC addresses. When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the switch.

To disable MAC learning for a set of VLANs, include the **no-mac-learning** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options]
```

no-mac-learning;

Enabling MAC Accounting

By default, MAC accounting is disabled. You can enable packet accounting either for a router or switch as a whole or for a specific VLAN. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned.

To enable MAC accounting, include the **global-mac-statistics** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-mac-statistics;
```

Enabling MAC Accounting for a VLAN

By default, MAC accounting is disabled. You can enable packet counting for a VLAN. When you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the interfaces in the VLAN.

To enable MAC accounting for a VLAN, include the **mac-statistics** statement at the **[edit vlans *vlan-name* switch-options]** hierarchy level:

```
[edit vlans vlan-name switch-options]
mac-statistics;
```

Enabling MAC Accounting for a Set of VLANs

By default, MAC accounting is disabled. You can enable packet counting for a set of VLANs. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the trunk port associated with the set of VLANs.

To enable MAC accounting for a set of VLANs, include the **mac-statistics** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options on page 62]
mac-statistics;
```

Configuring Inner and Outer TPIDs and VLAN IDs

For some rewrite operations, you must configure the inner or outer TPID values and inner or outer VLAN ID values. These values can be applied to either the input VLAN map or the output VLAN map.

On Ethernet IQ, IQ2, and IQ2-E interfaces; on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces; and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to configure the inner TPID, include the **inner-tag-protocol-id** statement:

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

You can include this statement at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number* input-vlan-map**]**
- **[edit interfaces** *interface-name* **unit** *logical-unit-number* output-vlan-map**]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* input-vlan-map**]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* output-vlan-map**]**

For the inner VLAN ID, include the **inner-vlan-id** statement. For the outer TPID, include the **tag-protocol-id** statement. For the outer VLAN ID, include the **vlan-id** statement:

```
input-vlan-map {
   (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
   inner-tag-protocol-id tpid;
   inner-vlan-id number;
   tag-protocol-id tpid;
   vlan-id number;
}
output-vlan-map {
   (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
   inner-tag-protocol-id tpid;
   inner-vlan-id number;
   tag-protocol-id tpid;
   vlan-id number;
}
```

For aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, include the **tag-protocol-id** statement for the outer TPID. For the outer VLAN ID, include the **vlan-id** statement:

```
input-vlan-map {
   (pop | push | swap);
   tag-protocol-id tpid;
   vlan-id number;
}
output-vlan-map {
   (pop | push | swap);
   tag-protocol-id tpid;
   vlan-id number;
}
```

You can include these statements at the following hierarchy levels:

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

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see 802.1Q VLANs Overview.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* ether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

Table 5 on page 29 and Table 6 on page 29 specify when these statements are required. Table 5 on page 29 indicates valid statement combinations for rewrite operations for the input VLAN map. "No" means the statement must not be included in the input VLAN map for the rewrite operation. "Optional" means the statement may be optionally specified for the rewrite operation in the input VLAN map. "Any" means that you must include the **vlan-id** statement, **tag-protocol-id** statement, **inner-vlan-id** statement, or **inner-tag-protocol-id** statement.

Table 5: Rewrite Operations and Statement Usage for Input VLAN Maps

	Input VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Optional	Optional	No	No
pop	No	No	No	No
swap	Any	Any	No	No
push-push	Optional	Optional	Optional	optional
swap-push	Optional	Optional	Any	Any
swap-swap	Optional	Optional	Any	Any
pop-swap	No	No	Any	Any
pop-pop	No	No	No	No

Table 6 on page 29 indicates valid statement combinations for rewrite operations for the output VLAN map. "No" means the statement must not be included in the output VLAN map for the rewrite operation. "Optional" means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 6: Rewrite Operations and Statement Usage for Output VLAN Maps

	Output VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	No	Optional	No	No
pop	No	No	No	No
swap	No	Optional	No	No
push-push	No	Optional	No	Optional

Table 6: Rewrite Operations and Statement Usage for Output VLAN Maps *(continued)*

	Output VLAN Map Statements			
swap-push	No	Optional	No	Optional
swap-swap	No	Optional	No	Optional
pop-swap	No	No	No	Optional
pop-pop	No	No	No	No

The following examples use Table 5 on page 29 and Table 6 on page 29 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

Input VLAN Map with inner-vlan-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
  inner-tag-protocol-id tpid;
}
```

Stacking a VLAN Tag

To stack a VLAN tag on all tagged frames entering or exiting the interface, include the **push**, **vlan-id**, and **tag-protocol-id** statements in the input VLAN map or the output VLAN map:

```
input-vlan-mapinput-vlan-map {
    push;
    vlan-id number;
    tag-protocol-id tpid;
}
output-vlan-map {
    push;
    tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

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

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces** *interface-name* **ether-options ethernet-switch-profile tag-protocol-id [** *tpids* **]]** hierarchy level.

Configuring the Size of the MAC Address Table

You can modify the size of the MAC address table for each VLAN. The default table size is 5120 addresses. The minimum you can configure is 16 addresses, and the maximum is 1,048,575 addresses.

If the MAC table limit is reached, new addresses can no longer be added to the table. Unused MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added.

To modify the size of the MAC table, include the **mac-table-size** *limit* statement at the **[edit vlans** *vlan-name* **switch-options]** hierarchy level:

```
[edit]
vlans {
    vlan-name {
        domain-type bridge;
        switch-options {
```

```
        mac-table-size limit {
          packet-action drop;
        }
      }
    }
  }
```

Related Documentation

- Disabling MAC Learning for a Bridge Domain or Logical Interface
- Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain
- Limiting MAC Addresses Learned from an Interface in a Bridge Domain
- Enabling MAC Accounting for a Bridge Domain

Configuring Static MAC Addresses for Logical Interfaces in a VLAN

You can manually add static MAC entries for the logical interfaces in a VLAN. You can specify one or more static MAC addresses for each logical interface.

To add a static MAC address for a logical interface in a VLAN, include the **static-mac** *mac-address* statement at the **[edit vlans** *vlan-name* **switch-options interface** *interface-name***]** hierarchy level.

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    switch-options {
      interface interface-name {
        static-mac mac-address {
          <vlan-id number>;
        }
      }
    }
  }
}
```

You can optionally specify a VLAN identifier for the static MAC address by using the **vlan-id** statement. To specify a VLAN identifier for a static MAC address, you must use the **all** option when configuring a VLAN identifier for the VLAN.

NOTE: If a static MAC address you configure for a logical interface appears on a different logical interface, packets sent to that interface are dropped.

Related Documentation

- Disabling MAC Learning for a VLAN or Logical Interface on page 25
- Configuring the Size of the MAC Address Table on page 31
- Enabling MAC Accounting for a VLAN on page 27

Disabling Layer 2 Learning and Forwarding

Disabling dynamic MAC learning on an MX Series router or an EX Series switch prevents all the logical interfaces on the router from learning source and destination MAC addresses.

To disable MAC learning for an MX Series router or an EX Series switch, include the **global-no-mac-learning** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-no-mac-learning;
```

For information about how to configure a virtual switch, see Configuring a Layer 2 Virtual Switch.

Related Documentation

- Layer 2 Learning and Forwarding Overview
- Configuring the MAC Table Timeout Interval
- Enabling MAC Accounting
- Limiting the Number of MAC Addresses Learned from Each Logical Interface

Configuring Multiple VLAN Registration Protocol (MVRP)

Multiple VLAN Registration Protocol (MVRP) is used to manage dynamic VLAN registration in Carrier Ethernet network. You can use MVRP on MX Series routers or on EX Series switches.

For information about using MVRP on EX Series switches, see Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches.

MVRP is disabled by default on MX Series routers and EX Series switches.

To enable MVRP or set MVRP options, follow these instructions:

- Enabling MVRP on page 33
- Disabling MVRP on page 34
- Changing the Registration Mode to Disable Dynamic VLANs on page 34
- Configuring Timer Values on page 34
- Configuring the Multicast MAC address for MVRP on page 35
- Configuring an MVRP Interface as a Point-to-Point Interface on page 35
- Configuring MVRP Tracing Options on page 35

Enabling MVRP

MVRP can only be enabled on trunk interfaces.

To enable MVRP on a specific trunk interface (here, interface **ge-3/0/5**):

```
[edit protocols mvrp]
```

```
user@host# set interfaces ge-3/0/5
```

Disabling MVRP

MVRP is disabled by default. You only need to perform this procedure if you have previously enabled MVRP.

To disable MVRP on all trunk interfaces, use one of the following:

```
[edit protocols mvrp]
user@host# deactivate protocols mvrp
user@host# delete protocols mvrp
```

Changing the Registration Mode to Disable Dynamic VLANs

When the registration mode for an interface is set to **normal** (the default), dynamic VLANs are created on interfaces participating in MVRP. The dynamic VLANs created on one router device are then propagated by means of MVRP to other router devices in a topology.

However, Dynamic VLAN creation through MVRP can be disabled for all trunk interfaces on a router device or for individual trunk interfaces.

For information about disabling dynamic VLAN creation on an interface so that the interface does not register and does not participate in MVRP, see Controlling the Management State of a VLAN in MVRP Configurations (CLI Procedure).

Configuring Timer Values

The timers in MVRP define the amount of time an interface waits to join or leave MVRP or to send or process the MVRP information for the router or switch after receiving an MVRP PDU:

- The join timer controls the amount of time the router waits to accept a registration request.
- The leave timer controls the period of time that the router waits in the Leave state before changing to the unregistered state.
- The leaveall timer controls the frequency with which the LeaveAll messages are communicated.

The default MVRP timer values are 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer.

BEST PRACTICE: **Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.**

To set the join timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interfaces ge-3/0/5 join-timer 300
```

To set the leave timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interfaces ge-3/0/5 leave-timer 1200
```

To set the leaveall timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5 leaveall-timer 12000
```

Configuring the Multicast MAC address for MVRP

MVRP uses the customer MVRP multicast MAC address when MVRP is enabled. However, you can configure MVRP to instead use the provider MVRP multicast MAC address.

To configure MVRP to use the provider MVRP multicast MAC address:

```
[edit protocols mvrp]
user@host# set bpdu-destination-mac-address provider-bridge-group;
```

Configuring an MVRP Interface as a Point-to-Point Interface

Specify that a configured interface is connected point-to-point. If specified, a point-to-point subset of the MRP state machine provides a simpler and more efficient method to accelerate convergence on the network.

To specify that an MVRP interface is point-to-point (here, interface **ge-3/0/5**):

```
[edit protocols mvrp]
user@host# set interfaces ge-3/0/5 point-to-point;
```

Configuring MVRP Tracing Options

Set MVRP protocol-level tracing options.

To specify MVRP protocol tracing (here, the file is **/var/log/mvrp-log**, size is **2m**, number of files is **28**, the option **world-readable** indicates the log can be read by user, and MVRP is flagging **events**):

```
[edit protocols mvrp]
user@host# edit protocols mvrp traceoptions file /var/log/mvrp-log size 2m files 28
  world-readable flag events
```

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Verifying That MVRP Is Working Correctly

Configuring a Layer 2 Virtual Switch

A Layer 2 virtual switch, which isolates a LAN segment with its spanning-tree protocol instance and separates its VLAN ID space, filters and forwards traffic only at the data link layer. Each VLAN consists of a set of logical ports that participate in Layer 2 learning and forwarding. A virtual switch represents a Layer 2 network.

Two main types of interfaces are used in virtual switch hierarchies:

- Layer 2 logical interface—This type of interface uses the VLAN-ID as a virtual circuit identifier and the scope of the VLAN-ID is local to the interface port. This type of interface is often used in service-provider-centric applications.
- Access or trunk interface—This type of interface uses a VLAN-ID with global significance. The access or trunk interface is implicitly associated with VLANs based on VLAN membership. Access or trunk interfaces are typically used in enterprise-centric applications.

NOTE: The difference between access interfaces and trunk interfaces is that access interfaces can be part of one VLAN only and the interface is normally attached to an end-user device (packets are implicitly associated with the configured VLAN). In contrast, trunk interfaces multiplex traffic from multiple VLANs and usually interconnect switches.

To configure a Layer 2 virtual switch, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name (
    instance-type virtual-switch;
    vlans vlan-name{
      vlan-id (all | none | number);
        [...configure optional VLAN parameters]
    }
  }
}
```

To enable a virtual switch, you must specify **virtual-switch** as the **instance-type**.

The VLANs that are specified with the **vlan-id** statement are included in the virtual switch.

You can configure other optional VLAN parameters in the virtual switch.

Related Documentation

- Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 36

Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port

You can associate one or more Layer 2 trunk interfaces with a virtual switch.

A virtual switch configured with a Layer 2 trunk port also supports IRB within a VLAN. IRB provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. Only an interface configured with the **interface-mode (access | trunk)** statement can be associated with a virtual switch. An access interface enables you to accept packets with no VLAN identifier.

In addition, you can configure Layer 2 learning and forwarding properties for the virtual switch.

To configure a virtual switch with a Layer 2 trunk interface, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name {
    instance-type virtual-switch;
    interface interface-name;
    vlans name{
      vlan-id (all | none | number);
        [...configure optional VLAN parameters]
    }
  }
}
```

Related Documentation

- Configuring a Layer 2 Virtual Switch on page 35

Configuring VLAN Encapsulation

To configure encapsulation on an interface, enter the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation type;
```

The following list contains important notes regarding encapsulation:

- Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.
- For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.
- For some encapsulation types, including flexible Ethernet services, Ethernet VLAN CCC, and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement:

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

  You can include this statement at the following hierarchy levels:

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

- You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

In general, you configure an interface's encapsulation at the **[edit interfaces *interface-name*]** hierarchy level.

Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

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

Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

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

Related Documentation

- 802.1Q VLANs Overview
- Junos® OS Ethernet Interfaces

Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

```
proxy-arp (restricted |unrestricted);
```

You can include this statement at the following hierarchy levels:

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

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.

NOTE: When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See Configuring Gratuitous ARP for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- proxy-arp on page 107
- Restricted and Unrestricted Proxy ARP Overview on page 11
- Configuring Gratuitous ARP
- Junos® OS Ethernet Interfaces

Rewriting a VLAN Tag and Adding a New Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and on Gigabit Ethernet and 10-Gigabit Ethernet interfaces on EX Series switches, to replace the outer VLAN tag of the incoming frame with a user-specified VLAN tag value, include the **swap-push** statement in the input VLAN map or output VLAN map:

swap-push

A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

See Configuring Inner and Outer TPIDs and VLAN IDs and Configuring Inner and Outer TPIDs and VLAN IDs for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 82
- output-vlan-map on page 100
- swap-push on page 115

- unit
- Junos® OS Ethernet Interfaces

Configuring VLAN Translation with a VLAN ID List

In many cases, the VLAN identifiers on the frames of an interface's packets are not correct. VLAN translation, or VLAN rewrite, allows you to configure bidirectional VLAN identifier translation with a list on frames arriving on and leaving from a logical interface. This lets you use unique VLAN identifiers internally and maintain legacy VLAN identifiers on logical interfaces.

To perform VLAN translation on the packets on a trunk interface, insert the **vlan-rewrite** statement at the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level. You must also include the **interface-mode trunk** statement within the **[edit interfaces *interface-name* unit *unit-number* family ethernet-switching]** hierarchy because VLAN translation is only supported on trunk interfaces. The reverse translation takes place on traffic exiting the interface. In other words, if VLAN 200 is translated to 500 on traffic entering the interface, VLAN 500 is translated to VLAN 200 on traffic leaving the interface.

The following example translates incoming trunk packets from VLAN identifier 200 to 500 and 201 to 501 (other valid VLAN identifiers are not affected):

```
[edit interfaces ge-1/0/1]
unit 0 {
   ... # Other logical interface statements
   family ethernet-switching {
      interface-mode trunk # Translation is only for trunks
      inner-vlan-id-list [ 100 500–600 ];
      vlan-rewrite {
         translate 200 500;
         translate 201 501;
      }
   ... # Other ethernet-switching statements
   }
}
```

NOTE: This example also translates frame VLANs from 500 to 200 and 501 to 201 on egress.

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag

Configuring a Logical Interface for Access Mode

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified VLAN. A logical interface configured to accept untagged packets is called an *access interface* or *access port*.

interface-mode access;

You can include this statement at the following hierarchy levels:

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

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

The following example configures a logical interface as an access port with a VLAN ID of 20 on routers and switches that support the enhanced Layer 2 software:

```
[edit interfaces ge-1/2/0]
unit 1 {
   family ethernet-switching {
      interface-mode access;
      vlan members 20;
   }
}
```

Related Documentation

- 802.1Q VLANs Overview
- Junos® OS Ethernet Interfaces

CHAPTER 5

Configuration Statements

[edit dynamic-profiles] Hierarchy Level

```
dynamic-profiles {
  profile-name {
    class-of-service
    ... statements from those in [edit class-of-service] Hierarchy Level.
    firewall
    ... statements from those in [edit firewall] Hierarchy Level.
    interfaces
    ... statements from those in [edit interfaces] Hierarchy Level.
    policy-options
    ... statements from those in [edit policy-options] Hierarchy Level.
    predefined-variable-defaults variable-name default-value
    profile-variable-set variable-set-name dynamic-variable-name substitute-variable-name
    protocols
    ... statements from those in [edit protocols] Hierarchy Level.
    routing-instances
    ... statements from those in [edit routing-instances] Hierarchy Level.
    routing-options
    ... statements from those in [edit routing-options] Hierarchy Level.
    services
    ... statements from those in [edit services] Hierarchy Level.
      variables {
        variable-name {
          default-value default-value;
          equals expression;
          mandatory;
          }
          uid;
          uid-reference;
```

```
        }
      }
    }
```

Related Documentation

- Notational Conventions Used in Junos OS Configuration Hierarchies
- [edit dynamic-profiles routing-instances] Hierarchy Level
- [edit dynamic-profiles routing-options] Hierarchy Level
- [edit dynamic-profiles variables] Hierarchy Level

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy
      level ...
  }
  interface-set interface-set-name {
    interface interface-name {
      (unit unit-number | vlan-tags-outer vlan-tag);
    }
  }
  irb (Interfaces) {
    accounting-profile name;
    description  text;
    disable;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      bandwidth rate;
      description text;
      disable;
      encapsulation type;
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            output;
          }
```

```
    }
    address ipv4-address {
      arp ip-address (mac | multicast-mac) mac-address <publish>;
      broadcast address;
      preferred;
      primary;
      vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
          hold-time seconds;
        }
        priority number;
        track {
          interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
          }
          priority-hold-time seconds;
          route prefix/prefix-length routing-instance instance-name priority-cost priority;
        }
        virtual-address [ addresses ];
        vrrp-inherit-from vrrp-group;
      }
    }
    filter {
      input filter-name;
      output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    primary;
    rpf-check {
      fail-filter filter-name;
      mode {
        loose;
      }
    }
    targeted-broadcast {
      forward-and-send-to-re;
      forward-only;
    }
  }
  family inet6 {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        output;
      }
    }
```

```
    address address {
      eui-64;
      ndp ip-address (mac | multicast-mac) mac-address <publish>;
      preferred;
      primary;
      vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
          hold-time seconds;
        }
        priority number;
        track {
          interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
          }
          priority-hold-time seconds;
          route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6–address;
        vrrp-inherit-from {
          active-group group-number;
          active-interface interface-name;
        }
      }
    }
    (dad-disable | no-dad-disable);
    filter {
      input filter-name;
      output filter-name;
    }
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    no-redirects;
    policer {
      input policer-name;
      output policer-name;
    }
    rpf-check {
      fail-filter filter-name;
      mode {
        loose;
      }
    }
  }
  family iso {
    address interface-address;
    mtu bytes;
  }
```

```
        family mpls {
          filter {
            input filter-name;
            output filter-name;
          }
          mtu bytes;
          policer {
            input policer-name;
            output policer-name;
          }
        }
        native-inner-vlan-id vlan-id;
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id-list [vlan-id's];
        vlan-id-range [vlan-id-range];
      }
    }
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      flag flag <disable>;
      no-remote-trace;
    }
}

interfaces {
  interface-name {
    disable;
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id [ hexadecimal-identifiers ];
      }
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        admin-key key;
        fast-failover;
        link-protection {
          disable;
          (revertive | non-revertive);
        }
        periodic (fast | slow);
        system-id mac-address;
        system-priority priority;
      }
      (link-protection | no-link-protection);
      link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
      logical-interface-fpc-redundancy;
      (loopback | no-loopback);
      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);
      }
      minimum-links number;
      rebalance-periodic {
        start-time time;
        interval number;
      }
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
    auto-configure {
      remove-when-no-subscribers;
      stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
          password password-string;
          username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82 ( circuit-id | remote-id);
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
          }
        }
        dynamic-profile profile-name {
          accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
          ranges (any | low-tag–high-tag),(any | low-tag–high-tag);
        }
      }
      vlan-ranges {
        access-profile profile-name;
        authentication {
          password password-string;
          username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
          }
        }
        dynamic-profile profile-name {
```

```
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)–(any | high-tag);
      }
    }
    override tag vlan-tag dynamic-profile profile name;
  }
  encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
  ether-options {
    802.3ad {
      aex;
      (backup | primary);
      lacp {
        force-up;
        port-priority
      }
    }
    asynchronous-notification;
    (auto-negotiation | no-auto-negotiation);
    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;
          }
        }
        tag-protocol-id;
      }
      (mac-learn-enable | no-mac-learn-enable);
    }
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    link-mode (automatic | full-duplex | half-duplex);
    (lloopback | no-loopback);
    keepalives <interval seconds> <down-count number> <up-count number>;
    speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
    source-address-filter {
      mac-address;
    }
```

```
    source-filtering | no-source-filtering;
  }
  flexible-vlan-tagging;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time (up milliseconds | down milliseconds);
  interface-transmit-statistics;
  (keepalives <down-count number> <interval seconds> <up-count number> |
    no-keepalives);
  layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
  }
  link-mode (automatic | full-duplex);
  mac mac-address;
  mtu bytes;
  multi-chassis-protection peer-ip-address {
    interface interface-name;
  }
  native-vlan-id number;
  no-gratuitous-arp-request;
  optics-options {
    alarm low–light–alarm {
      (link-down | syslog);
    }
    warning low–light–warning {
      (link-down | syslog);
    }
    wavelength nm;
  }
  passive-monitor-mode;
  per-unit-scheduler;
  speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
  stacked-vlan-tagging;
  traceoptions {
    flag flag;
  }
  transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
  }
  (traps | no-traps);
  unidirectional;
  vlan-tagging;
}


interface-name {
  unit logical-unit-number {
    disable;
    accept-source-mac {
      mac-address mac-address {
        policer {
          input policer-name;
          output policer-name;
        }
```

```
    }
}
account-layer2-overhead (Interface Level) {
    value;
    egress bytes;
    ingress bytes;
}
accounting-profile name;
advisory-options {
    downstream-rate rate;
    upstream-rate rate;
}
arp-resp (restricted|unrestricted);
bandwidth rate;
clear-dont-fragment-bit;
copy-tos-to-outer-ip-header;
demux-destination family;
encapsulation (vlan-bridge | vlan-vpls);
epd-threshold cells plp1 cells;
filter filter-name;
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;
}
interface-shared-with psdnumerical-index;
layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
```

```
    vlan-id number;
    vlan-id-list [ vlan-id vlan-id-vlan-id ];
    vlan-id-range number-number;
    vlan-tags (inner <tpid.>vlan-id | inner-list [vlan-id vlan-id–vlan-id ] |
      inner-range <tpid.>vlan-id–vlan-id) outer <tpid.>vlan-id;
  }

  unit logical-unit-number {
    family ethernet-switching {
      filter{
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
      (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
        number–number ]);
      interface-mode (access | trunk);
      policer {
        input policer-name;
        output policer-name;
      }
      vlan-rewrite {
        translate old-vlan-id new-vlan-id;
      }
      vlan {
        members [ all vlan-identifiers ];
      }
    }
    family inet {
      filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
      }
      input-hierarchical-policer policer-name;
      mac-validate (loose | strict);
      mtu bytes;
      no-neighbor-learn;
      no-redirects;
      policer {
        arp policer-template-name;
        input policer-name;
        output policer-name;
      }
      primary;
      receive-options-packets;
      receive-ttl-exceeded;
      rpf-check {
        fail-filter filter-name;
        mode loose;
      }
      sampling {
        (input | output | input output);
      }
      simple-filter {
        input filter-name;
      }
```

```
    targeted-broadcast {
      forward-and-send-to-re;
      forward-only;
    }
    unnumbered-address interface-name <destination address>
      <destination-profile profile-name> <preferred-source-address address>;

  }

  family inet6 {
    address ipv6-address {
      destination destination-address;
      eui-64;
      ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
        multicast-mac multicast-mac-address) <publish>>;
      preferred;
      primary;
      vrrp-inet6-group group-number {
        (accept-data | no-accept-data);
        fast-interval milliseconds;
        inet6-advertise-interval seconds;
        (no-preempt; | ... the following preempt statement ...)
        preempt {
          hold-time seconds;
        }
        priority number;
        track {
          interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
          }
          priority-hold-time seconds;
          route ip-address-prefix/prefix-length routing-instance instance-name
            priority-cost priority;
        }
        virtual-inet6-address [ addresses ];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
          active-group group-number;
          active-interface interface-name;
        }
      }
    }
    (dad-disable | no-dad-disable);
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
      input policer-name;
      output policer-name;
```

```
            }
            rpf-check {
              fail-filter filter-name;
              mode loose;
            }
            sampling {
              (input | output | input output);
            }
            unnumbered-address interface-name preferred-source-address address;

          }

          family iso {
            address iso-address;
            mtu bytes;
          }

          family mlfr-end-to-end {
            bundle logical-interface-name;
          }



          family mpls {
            filter {
              group filter-group-number;
              (input filter-name | input-list [ filter-names ]);
              (output filter-name | output-list [ filter-names ]);
            }
            input-hierarchical-policer policer-name;
            maximum-labels maximum-labels;
            mtu bytes;
            policer {
              input policer-name;
              output policer-name;
            }
          }



          family vpls {
            core-facing;
            filter {
              group filter-group-number;
              (input filter-name | input-list [ filter-names ]);
              (output filter-name | output-list [ filter-names ]);
            }
            policer {
              input policer-name;
              output policer-name;
            }
          }

        }
      }
    }
```

Related Documentation
- Notational Conventions Used in Junos OS Configuration Hierarchies

[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 Junos OS Routing Protocols Configuration Guide.

```
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*
- Junos® OS Ethernet Interfaces
- Junos® OS Network Interfaces

[edit protocols l2-learning] Hierarchy Level

```
protocols {
  l2-learning {
    global-mac-limit {
      limit;
      packet-action drop;
    }
    global-mac-statistics;
    global-mac-table-aging-time seconds;
    global-no-mac-learning;
  }
}
```

Related Documentation

- Notational Conventions Used in Junos OS Configuration Hierarchies
- [edit protocols] Hierarchy Level

Layer 2 Routing Instances Configuration Hierarchy

Use the **vpls** routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN.

To configure routing instances for Layer 2 networks, include the following statements:

```
routing-instances {
  routing-instance-name {
    access {
      address-assignment {
        ... same statements as in the address-assignment subhierarchy in [edit access]
          Hierarchy Level ...
      }
      address-protection;
    description text;
    egress-protection {
      context-identifier context-id;
    }
    forwarding-options {
      ...forwarding-options...
    }
    instance-role role;
    instance-type type;
    interface interface-name;
    l2-domain-id-for-l3 id;
    l2vpn-id community;
    layer3-domain-identifier identifier;
    multicast-snooping-options {
      ... same statements as in [edit multicast-snooping-options] Hierarchy Level EXCEPT
        FOR ...
      traceoptions {...}   # NOT valid at this level
    }
    no-irb-layer-2-copy;
    no-local-switching;
    no-vrf-advertise;
    no-vrf-propagate-ttl;
    pbb-options {
      default-bvlan bvlan;
      peer-instance instance;
      vlan-id vlan-id isid-list [ isid-numbers ]
    }
    protocols {
      ... the protocols subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
    }
    provider-tunnel {
      ... the provider-tunnel subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
    }
```

```
route-distinguisher (as-number:number | ip-address:number);
routing-interface interface;
routing-options {
  ... the routing-options subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
service-groups {
  service-group-name {
    pbb-service-options {
      default-isid isid-number;
      isid isid-number vlan-id-list [ vlan-ids ];
      mac-address mac-address;
    }
    service-type type;
  }
}
services {
  mobile-ip {
    ... same statements as in [edit services mobile-ip] Hierarchy Level ...
  }
}
switch-options {
  ... same statements as in [edit switch-options] Hierarchy Level ...
}
vlan-id (id | all | none);
vlan-model one-to-one;
vlan-tags outer <tpid.>vlan-id inner <tpid.>vlan-id;
[edit vlans] Hierarchy Level on page 63 {
  ... same statements as in [edit vlans] Hierarchy Level ...
}
vrf-advertise-selective {
  family {
    inet-mvpn;
    inet6-mvpn;
  }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-propagate-ttl;
vrf-table-label;
vrf-target {
  export community-name;
  import community-name;
}
protocols {
  ... protocols-configuration ...
}
routing-options {
  ... routing-options-configuration ...
}
bridge-domains {
  bridge-domain-name {
    domain-type bridge;
    interface interface-name;
    routing-interface routing-interface-name;
    vlan-id (Bridge Domain or VLAN) (none | all | number);
```

```
            vlan-tags outer number inner number;
            bridge-options {
              interface-mac-limit limit {
                packet-action drop;
              }
              interface interface-name {
                interface-mac-limit limit {
                  packet-action drop;
                }
              }
              mac-statistics;
              mac-table-size limit {
                packet-action drop;
              }
              no-mac-learning;
              static-mac mac-address;
            }
          }
        }
      }
    }
```

With the exception of the **instance-type virtual-switch** statement (which configures a virtual-switch routing instance), you can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

Related Documentation

- Routing Instances Overview
- Layer 2 Routing Instance Types
- Configuring a Layer 2 Virtual Switch on page 35
- Configuring a Layer 2 Control Protocol Routing Instance

[edit switch-options] Hierarchy Level

```
switch-options {
  interface interface-name {
    interface-mac-limit {
      number-of-addresses;
      packet-action drop;
    }
    no-mac-learning;
  }
  interface-mac-limit {
    number-of-addresses;
    packet-action drop;
  }
```

```
    mac-statistics;
    mac-table-size {
      number-of-addresses;
      packet-action drop;
    }
    no-mac-learning;
    service-id number;
  }
```

[edit vlans] Hierarchy Level

```
vlans {
  vlan-name {
    description text-description;
    domain-type bridge;
    forwarding-options {
      filter {
        input filter-name;
      }
      flood {
        input filter-name;
      }
    }
    interface interface-name;
    l3-interface interface-name;
    multicast-snooping-options {
      ... same statements as in multicast-snooping-options ...
    }
    no-irb-layer-2-copy;
    service-id number;
    switch-options {
      ... the switch-options subhierarchy appears after the main [edit vlans vlan-name]
        hierarchy ...
    }
    vlan-id (all | none | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer <tpid.>vlan-id <inner <tpid.>vlan-id>;
  }

  vlan-name {
    switch-options {
      interface interface-name {
        interface-mac-limit {
          limit;
          packet-action drop;
        }
        no-mac-learning;
        static-mac mac-address {
          vlan-id number;
        }
      }
      interface-mac-limit {
        limit;
        packet-action drop;
      }
```

```
            mac-statistics;
            mac-table-size {
              number-of-addresses;
              packet-action drop;
            }
            no-mac-learning;
          }
        }

      }
```

bpdu-destination-mac-address

Syntax bpdu-destination-mac-address provider-bridge-group;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mvrp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit protocols mvrp],
[edit routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type)

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description For Multiple VLAN Registration Protocol (MVRP) configurations, specifies the multicast address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the Junos OS uses the customer MVRP multicast MAC address.

Default By default, the provider MVRP multicast MAC address is used (if configured). Otherwise, the customer MVRP MAC address is used.

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7

bridge-priority

Syntax bridge-priority *priority*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* protocols mstp msti *msti-id*],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mstp msti *msti-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*],
[edit protocols (mstp | rstp)],
[edit protocols mstp msti *msti-id*],
[edit protocols vstp vlan *vlan-id*],
[edit routing-instances *routing-instance-name* protocols (mstp | rstp)],
[edit routing-instances *routing-instance-name* protocols mstp msti *msti-id*],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*]

Release Information Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.

Description Determine which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.

Options ***priority***—The bridge priority can be set only in increments of 4096.

Range: 0 through 61,440

Default: 32,768

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

Related Documentation
- Bridge Priority for Election of Root Bridge and Designated Bridge

domain-type

Syntax domain-type bridge;

Hierarchy Level [edit bridge-domains *bridge-domain-name*],
[edit vlans on page 63 *vlan-name*]
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Define the type of domain for a Layer 2 bridge domain or VLAN.

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

Related Documentation

- Configuring a Bridge Domain
- Configuring a VLAN on page 18
- Configuring a Layer 2 Virtual Switch
- Configuring a Layer 2 Virtual Switch on page 35

encapsulation (Physical Interface)

Syntax encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc | ethernet-bridge | ethernet-ccc | ethernet-over-atm | ethernet-tcc | ethernet-vpls | ethernet-vpls-fr | ether-vpls-over-atm-llc | ethernet-vpls-ppp |extended-frame-relay-ccc | extended-frame-relay-ether-type-tcc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services | flexible-frame-relay | frame-relay | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-port-ccc | frame-relay-tcc | generic-services | multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc | vlan-vci-ccc | vlan-vpls);

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces rlsq *number:number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches (**flexible-ethernet-services**, **ethernet-ccc**, and **ethernet-tcc** options only).

Description Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Default **ppp**—Use serial PPP encapsulation.

Options **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.

atm-pvc—Use ATM PVC encapsulation.

cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.

cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.

cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.

ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.

ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination

IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

NOTE: **The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.**

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

NOTE: **Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.**

Required Privilege Level

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

Related Documentation

- Configuring Interface Encapsulation on Physical Interfaces
- Configuring CCC Encapsulation for Layer 2 VPNs
- Configuring Layer 2 Switching Cross-Connects Using CCC
- Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits
- Configuring ATM Interface Encapsulation
- Configuring ATM-to-Ethernet Interworking
- Configuring VLAN Encapsulation
- Configuring Extended VLAN Encapsulation
- Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces
- Configuring Interfaces for Layer 2 Circuits
- Configuring Interface Encapsulation on PTX Series Packet Transport Switches
- Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)
- Configuring MPLS LSP Tunnel Cross-Connects Using CCC
- Configuring TCC
- Configuring VPLS Interface Encapsulation
- Configuring Interfaces for VPLS Routing
- Defining the Encapsulation for Switching Cross-Connects
- Understanding Encapsulation on an Interface

family

Syntax

```
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      (input | output | input output);
    }
  }
  access-concentrator name;
  address address {
    ... the address subhierarchy appears after the main [edit interfaces interface-name unit
      logical-unit-number family family-name] hierarchy ...
  }
  bridge-domain-type (bvlan | svlan);
  bundle interface-name;
  core-facing;
  demux-destination {
    destination-prefix;
  }
  demux-source {
    source-prefix;
  }
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
  }
  interface-mode (access | trunk);
  ipsec-sa sa-name;
  isid-list all-service-groups;
  keep-address-and-control;
  mac-validate (loose | strict);
  max-sessions number;
  max-sessions-vsa-ignore;
  mtu bytes;
  multicast-only;
  negotiate-address;
  no-redirects:
  policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
  }
  primary;
  protocols [inet iso mpls];
  proxy inet-address address;
  receive-options-packets;
  receive-ttl-exceeded;
  remote (inet-address address | mac-address address);
```

```
rpf-check {
  fail-filter filter-name
  mode loose;
}
sampling {
  input;
  output;
}
service {
  input {
    post-service-filter filter-name;
    service-set service-set-name <service-filter filter-name>;
  }
  output {
    service-set service-set-name <service-filter filter-name>;
  }
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
  maximum-seconds>;
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address destination-profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  master-only;
  multipoint-destination address dlci dlci-identifier;
  multipoint-destination address {
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
        sustained rate);
      queue-length number;
    }
    vci vpi-identifier.vci-identifier;
  }
  preferred;
  primary;
  vrrp-groupgroup-id {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-key key;
    authentication-type authentication;
```

```
            fast-interval milliseconds;
            (preempt | no-preempt) {
              hold-time seconds;
            }
            priority number;
            track {
              interface interface-name {
                bandwidth-threshold bits-per-second priority-cost priority;
                priority-cost priority;
              }
              priority-hold-time seconds;
              route prefix routing-instance instance-name priority-cost priority;
              }
            }
            virtual-address [ addresses ];
            }
          virtual-link-local-addressipv6-address;
        }
      }
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Option **max-sessions-vsa-ignore** introduced in Junos OS Release 11.4.

Description Configure protocol family information for the logical interface.

NOTE: Not all subordinate stanzas are available to every protocol family.

Options *family*—Protocol family:

- **any**—Protocol-independent family used for Layer 2 packet filtering

NOTE: This option is not supported on T4000 Type 5 FPCs.

- **ethernet-switching**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ccc**—Circuit cross-connect protocol suite
- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—Multiprotocol Label Switching (MPLS)
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpls**—(M Series and T Series routers only) Virtual private LAN service

The remaining statements are explained separately.

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

Related Documentation

- Configuring the Protocol Family
- Example: Configuring E-LINE and E-LAN Services for a PBB Network
- Junos Services Interfaces Configuration Release 11.2

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

global-mac-limit

Syntax

```
global-mac-limit limit {
  packet-action drop;
}
```

Hierarchy Level [edit protocols l2-learning]

Release Information Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.

Description (MX Series routers and EX Series switches only) Limit the number of media access control (MAC) addresses learned from the logical interfaces on the router.

Default 393,215 MAC addresses

Options *limit*—Number of MAC addresses that can be learned systemwide.

Range: 20 through 1,048,575

The remaining statement is explained separately in the "Summary of Bridge Domain Configuration Statements" chapter.

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

Related Documentation

- Limiting the Number of MAC Addresses Learned from Each Logical Interface

global-mac-move

Syntax

```
global-mac-move {
  notification-time seconds;
  threshold-count count;
  threshold-time seconds;
}
```

Hierarchy Level [edit protocols l2-learning]

Release Information Statement introduced in Junos OS Release 9.4.
Support for logical systems added in Junos OS Release 9.6.

Description (MX Series routers and EX Series switches only) Set parameters for media access control (MAC) address move reporting.

Default By default, MAC moves notify every second, with a threshold time of 1 second and a threshold count of 50.

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

Related Documentation
- Configuring MAC Move Parameters

global-mac-statistics

Syntax global-mac-statistics;

Hierarchy Level [edit protocols l2-learning]

Release Information Statement introduced in Junos OS Release 9.2.
Support for logical systems added in Junos OS Release 9.6.

Description (MX Series routers and EX Series switches only) Enable MAC accounting for the entire router or switch.

Default disabled

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

Related Documentation
- Enabling MAC Accounting

global-mac-table-aging-time

Syntax global-mac-table-aging-time *seconds*;

Hierarchy Level [edit protocols l2-learning]

Release Information Statement introduced in Junos OS Release 9.2.
Support for logical systems added in Junos OS Release 9.6.

Description (MX Series routers and EX Series switches only) Configure the timeout interval for entries in the MAC table.

Default 300 seconds

Options ***seconds***—Time elapsed before MAC table entries are timed out and entries are deleted from the table.

Range: 10 through 1 million

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

Related Documentation
- Configuring the MAC Table Timeout Interval

global-no-mac-learning

Syntax global-no-mac-learning;

Hierarchy Level [edit protocols l2-learning]

Release Information Statement introduced in Junos OS Release 9.2.
Support for logical systems added in Junos OS Release 9.6.

Description (MX Series routers and EX Series switches only) Disable MAC learning for the entire router or switch.

Default MAC learning is enabled.

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

Related Documentation
- Disabling Layer 2 Learning and Forwarding on page 33

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 Specify the VLAN ID to rewrite for the inner tag of the final packet.

On MX Series routers, you can use this command 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.

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

input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP)

Syntax

```
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements introduced in Junos OS Release 8.1.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet SFPP interfaces only, define the rewrite profile to be applied to incoming frames on this logical interface.

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Stacking a VLAN Tag
- output-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP) on page 100

interface

Syntax interface *interface-name*;

Hierarchy Level [edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit vlans on page 63 *vlan-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2.
In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description (MX Series routers and EX Series switches only) Specify the logical interfaces to include in the bridge domain, VLAN, VPLS instance, or virtual switch.

Options *interface-name*—Name of a logical interface. For more information about how to configure logical interfaces, see the Junos® OS Network Interfaces.

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

Related Documentation

- Configuring a Bridge Domain
- Configuring a VLAN on page 18
- Configuring a Layer 2 Virtual Switch
- Configuring a Layer 2 Virtual Switch on page 35

interface (MVRP)

Syntax

```
interface (all | interface-name) {
  join-timer milliseconds;
  leave-timer milliseconds;
  leaveall-timer milliseconds;
  point-to-point;
  registration (forbidden | normal | restricted);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mvrp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit protocols mvrp],
[edit routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type)

Release Information

Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description

Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).

Default

By default, MVRP is disabled.

Options

all—All interfaces on the router or switch.

interface-name—Names of interface to be configured for MVRP.

The remaining 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

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7

interface (Spanning Tree)

Syntax

```
interface interface-name {
  bpdu-timeout-action {
    alarm;
    block;
  }
  cost cost;
  edge;
  mode (p2p | shared);
  no-root-port;
  priority interface-priority;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* protocols vstp vlan *vlan-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*],
[edit protocols (mstp | rstp | vstp)],
[edit protocols vstp vlan *vlan-id*],
[edit routing-instances *routing-instance-name* protocols (mstp | rstp | vstp)],
[edit routing-instances *routing-instance-name* protocols vstp vlan *vlan-id*]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.

Description

Configure the interface to participate in the RSTP or MSTP instance.

Options

interface-name—Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.

The remaining 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

- Spanning-Tree Instance Interface

interface-mac-limit

Syntax

```
interface-mac-limit limit {
    packet-action drop;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* bridge-options],
[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name*],
[edit logical-systems *logical-system-name* switch-options],
[edit logical-systems *logical-system-name* switch-options interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit routing-instances *routing-instance-name* switch-options],
[edit routing-instances *routing-instance-name* switch-options interface *interface-name*],
[edit switch-options],
[edit switch-options on page 62],
[edit switch-options interface *interface-name*],
[edit switch-options on page 62 interface *interface-name*],
[edit vlans on page 63 *vlan-name* switch-options],
[edit vlans on page 63 *vlan-name* switch-options interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options], [edit switch-options interface *interface-name*], [edit vlans *vlan-name* switch-options], and [edit vlans *vlan-name* switch-options interface *interface-name*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers or EX Series switches only) Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.

Default

1024 MAC addresses for each logical interface.

Options

limit—Maximum number of MAC addresses learned from an interface.

Range: 1 through 131,071 MAC addresses per interface

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Layer 2 Learning and Forwarding for Bridge Domains Overview
- Layer 2 Learning and Forwarding for VLANs Overview on page 4
- Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 5

interface-mode

Syntax interface-mode (access | trunk);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit interfaces *interface-name* unit *logical-unit-number* family ethernet-switching]
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos Release 12.3R2 for EX Series switches.

Description Determines whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see Configuring a Trunk Interface on a Bridge Network.

Options **access**—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the **vlan-id** statement.

trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the **vlan-id-list** statement.

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

Related Documentation

- Configuring a Logical Interface for Access Mode
- Configuring a Logical Interface for Trunk Mode

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

isid

Syntax isid *isid-number* vlan-id-list [*vlan-ids*];

Hierarchy Level [edit routing-instances *routing-instance-name* service-groups *service-group-name* pbb-service-options]

Release Information Statement introduced in JUNOS Release 10.0.

Description For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the service identifier (I-SID) for the customer routing instance (I-component) service group.

Options ***isid***—Service identifier. Enter an I-SID in the range from **256** through **16777214**.

vlan-id-list [*vlan-ids*]—List of service VLANs (S-VLANs).

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

Related Documentation
- Example: Configuring E-LINE and E-LAN Services for a PBB Network

isid-list

Syntax isid-list all-service-groups;

Hierarchy Level [edit interfaces *pseudo-logical-interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in JUNOS Release 10.0.

Description For IEEE 802.1ah provider backbone bridge (PBB) configurations, map all service identifiers (I-SIDs) specified for the service groups.

Options **all-service-groups**—Map all service identifiers (I-SIDs) for the specified service groups.

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

Related Documentation

- Example: Configuring E-LINE and E-LAN Services for a PBB Network

join-timer (MVRP)

Syntax join-timer *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mvrp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type),
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit logical-systems *logical-system-name* protocols mvrp interface (all | *interface-name*)],
[edit protocols mvrp interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type)

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description For Multiple VLAN Registration Protocol (MVRP), configure the maximum interval interfaces must wait before sending MVRP protocol data units (PDUs).

Default 200 milliseconds

Options ***milliseconds***—Interval that the interface must wait before sending MVRP PDUs. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7
- leaveall-timer on page 93
- leave-timer on page 94
- point-to-point on page 103
- registration on page 110

l2-learning

Syntax

```
l2-learning {
  global-mac-limit limit;
  global-mac-statistics;
  global-mac-table-aging-time seconds;
  global-no-mac-learning;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 8.4.

Description (MX Series routers and EX Series switches only) Configure Layer 2 address learning and forwarding properties globally.

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

- Layer 2 Learning and Forwarding Overview

l3-interface

Syntax l3-interface *interface-name-logical-unit-number*;

Hierarchy Level [edit interfaces ge-*chassis/slot/port* unit *logical-unit-number* family ethernet-switching]
[edit vlans *vlan-name*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
[edit vlans *vlan-name*] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.

Description Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

Default No Layer 3 (routing) interface is associated with the VLAN.

Options *interface-name–logical-unit-number*—Name of a logical interface.

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

leaveall-timer (MVRP)

Syntax leaveall-timer *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mvrp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type),
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit logical-systems *logical-system-name* protocols mvrp interface (all | *interface-name*)],
[edit protocols mvrp interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type)

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.

Default 10000 milliseconds

Options ***milliseconds***—Interval between the sending of Leave All messages. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7
- join-timer on page 91
- leave-timer on page 94
- point-to-point on page 103
- registration on page 110

leave-timer (MVRP)

Syntax leave-timer *milliseconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mvrp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type),
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit logical-systems *logical-system-name* protocols mvrp interface (all | *interface-name*)],
[edit protocols mvrp interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type)

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.

Default 1000 milliseconds

Options
milliseconds—Interval that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the leave-timer interval at twice the join-timer interval. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7
- join-timer on page 91
- leaveall-timer on page 93
- point-to-point on page 103
- registration on page 110

mac-statistics

Syntax mac-statistics;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options],
[edit logical-systems *logical-system-name* switch-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit routing-instances *routing-instance-name* switch-options],
[edit switch-options],
[edit switch-options on page 62],
[edit vlans on page 63 *vlan-name* switch-options]

Release Information
Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options] and [edit vlans *vlan-name* switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description (MX Series routers and EX Series switches only) Enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.

Default disabled

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

Related Documentation

- Layer 2 Learning and Forwarding for Bridge Domains Overview
- Layer 2 Learning and Forwarding for VLANs Overview on page 4
- Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 5

mac-table-size

Syntax

```
mac-table-size limit {
  packet-action drop;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options],
[edit logical-systems *logical-system-name* switch-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit routing-instances *routing-instance-name* switch-options],
[edit switch-options],
[edit switch-options on page 62],
[edit vlans on page 63 *vlan-name* switch-options]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options] and [edit vlans *vlan-name* switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.

Options

limit—Specify the maximum number of addresses in the MAC address table.

Range: 16 through 1,048,575 MAC addresses

Default: 5120 MAC addresses

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Layer 2 Learning and Forwarding for Bridge Domains Overview
- Layer 2 Learning and Forwarding for VLANs Overview on page 4
- Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports

- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 5

mvrp

Syntax

```
mvrp {
  bpdu-destination-mac-address provider-bridge-group;
  join-timer milliseconds;
  leave-timer milliseconds;
  leaveall-timer milliseconds;
  interface (all | interface-name) {
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    point-to-point;
    registration (forbidden | normal | restricted);
  }
  no-dynamic-vlan;
  traceoptions (Spanning Trees) {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]
(for virtual switch instance type)
[edit protocols],
[edit routing-instances *routing-instance-name* protocols] (for virtual switch instance type),

Release Information

Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description

For Layer 2 networks, configure Multiple VLAN Registration Protocol (MVRP) to dynamically share VLAN information and dynamically configure needed VLANs. Maintaining VLAN configurations based on active VLANs reduces the amount of traffic traveling in the network, saving network resources. MVRP is configured on trunk interfaces.

The remaining statements are explained separately.

Default

MVRP is disabled by default.

Required Privilege Level

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7

native-vlan-id

Syntax native-vlan-id *number*;

Hierarchy Level [edit interfaces ge-*fpc/pic/port*]
[edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs configured for 802.1Q flexible VLAN tagging, for all Ethernet interfaces on MX Series routers, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, configure mixed tagging support for untagged packets on a port. When the **native-vlan-id** statement is included with the flexible-vlan-tagging statement, untagged packets are accepted on the same mixed VLAN-tagged port.

The logical interface on which untagged packets are received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

When the **native-vlan-id** statement is included with the interface-mode the statement, untagged packets are accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.

NOTE: As per 12.3 Ethernet Interfaces Configuration Guide, native-vlan-id enables the tagged interface to accept untagged frames. However, when native-vlan-id is configured, VLAN ID is popped in egress as long as the VID in the outgoing frame matches the native-vlan-id.

Options *number*—VLAN ID number.
Range: (ACX Series routers and EX Series switches) 0 through 4094.

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

Related Documentation

- Configuring Mixed Tagging Support for Untagged Packets
- Configuring a Logical Interface for Access Mode
- flexible-vlan-tagging

no-mac-learning

Syntax no-mac-learning;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options],
[edit logical-systems *logical-system-name* switch-options],
[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit routing-instances *routing-instance-name* switch-options],
[edit switch-options],
[edit switch-options on page 62],
[edit switch-options on page 62 interface *interface-name*],
[edit vlans on page 63 *vlan-name* switch-options],
[edit vlans on page 63 *vlan-name* switch-options interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options], [edit switch-options interface *interface-name*], [edit vlans *vlan-name* switch-options], and [edit vlans *vlan-name* switch-options interface *interface-name*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description (MX Series routers and EX Series switches only) Disable MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.

NOTE: When MAC learning is disabled for a VPLS routing instance, traffic is not load balanced and only one of the equal-cost next hops is used.

Default MAC learning is enabled. Use **no-mac-learning** to disable MAC learning.

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

Related Documentation
- Layer 2 Learning and Forwarding for Bridge Domains Overview
- Layer 2 Learning and Forwarding for VLANs Overview on page 4
- Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 5

output-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP)

Syntax
```
output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information
Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements added in Junos OS Release 8.1.

Description
For Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, define the rewrite operation to be applied to outgoing frames on this logical interface.

The statements are explained separately.

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

Related Documentation
- Stacking and Rewriting Gigabit Ethernet VLAN Tags
- input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP) on page 82

packet-action

Syntax packet-action drop;

Hierarchy Level
[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface-mac-limit *limit*],
[edit logical-systems *logical-system-name* switch-options interface-mac-limit *limit*],
[edit protocols l2-learning global-mac-limit *limit*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* interface-mac-limit *limit*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface-mac-limit *limit*],
[edit routing-instances *routing-instance-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit routing-instances *routing-instance-name* switch-options interface-mac-limit *limit*],
[edit switch-options interface-mac-limit *limit*],
[edit switch-options on page 62 interface *interface-name* interface-mac-limit *limit*],
[edit switch-options on page 62 interface-mac-limit *limit*],
[edit switch-options on page 62 mac-table-size *limit*],
[edit vlans on page 63 *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*],
[edit vlans on page 63 *vlan-name* switch-options interface-mac-limit *limit*],
[edit vlans on page 63 *vlan-name* switch-options mac-table-size *limit*]

Release Information
Statement introduced in Junos OS Release 8.4.
Support for the **switch-options** statement added in Junos OS Release 9.2.
Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
Support for logical systems added in Junos OS Release 9.6.
[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description (MX Series routers and EX Series switches only) Specify that packets for new source MAC addresses be dropped after the MAC address limit is reached. If this statement is not configured, then packets for new source MAC addresses are forwarded by default.

Default Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

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

Related Documentation

- Layer 2 Learning and Forwarding for Bridge Domains Overview
- Layer 2 Learning and Forwarding for VLANs Overview on page 4
- Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports
- Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 5

point-to-point (MVRP)

Syntax point-to-point;

Hierarchy Level [edit logical-systems *logical-system-name* protocols mvrp interface (all | *interface-name*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type),
[edit protocols mvrp interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type)

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description (Optional) For Multiple VLAN Registration Protocol (MVRP) configurations, configure an interface to be recognized as a point-to-point connection. If specified, a point-to-point subset of the MRP state machine is used to provide a simpler and more efficient method to accelerate convergence on the network.

Default By default, MVRP is disabled.

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7
- join-timer on page 91
- leaveall-timer on page 93
- leave-timer on page 94
- registration on page 110

pop

Syntax pop;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

On MX series routers, this statement can be used on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces.

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

Related Documentation

- Removing a VLAN Tag

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 Specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.

On MX Series routers, 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.,

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.

On MX Series routers, you can use this statement for 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.

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

proxy-arp

Syntax proxy-arp (restricted | unrestricted);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
restricted added in Junos OS Release 10.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for the QFX Series..

Description For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.

Default Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.

Options

- **none**—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
- **restricted**—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.
- **unrestricted**—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.

Default: unrestricted

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

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 38
- Configuring Proxy ARP (CLI Procedure)
- Example: Configuring Proxy ARP on an EX Series Switch
- Configuring Gratuitous ARP

push

Syntax push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.

On MX Series routers, You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces.

If you include the **push** statement in the configuration, you must also include the pop statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level.

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

Related Documentation
- Stacking a VLAN Tag

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.

On MX Series routers, You can use this command for 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. .

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

registration

Syntax registration (forbidden | normal | restricted);

Hierarchy Level [edit protocols mvrp interface (all | *interface-name*)];
[edit routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type);
[edit logical-systems *logical-system-name* protocols mvrp interface (all | *interface-name*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp interface (all | *interface-name*)] (for virtual switch instance type);

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description For Multiple VLAN Registration Protocol (MVRP) configurations, configure the registration mode for the interface.

Default normal

Options

forbidden—The interface or interfaces do not register and do not participate in MVRP.

normal—The interface or interfaces accept MVRP messages and participate in MVRP.

restricted—The interface or interfaces ignore all MVRP JOIN messages received for VLANs that are not statically configured for MVRP on the 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 Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7
- join-timer on page 91
- leaveall-timer on page 93
- leave-timer on page 94
- point-to-point on page 103

rstp

Syntax

```
rstp {
  bpdu-block-on-edge;
  bpdu-destination-mac-address provider-bridge-group;
  bridge-priority priority;
  extended-system-id;
  force-version stp;
  forward-delay seconds;
  hello-time seconds;
  max-age seconds;
  interface interface-name {
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost cost;
    edge;
    mode (p2p | shared);
    no-root-port;
    priority interface-priority;
  }
  priority-hold-time seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced in Junos OS Release 8.4.
bpdu-block-on-edge statement added in Junos OS Release 9.4.
bpdu-timeout-action statement added in Junos OS Release 9.4.
Support for logic systems added in Junos OS Release 9.6.

Description

Configure RSTP parameters.

Options

The statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring Rapid Spanning-Tree Protocol

service-id

Syntax service-id *number*;

Hierarchy Level [edit switch-options on page 62]
[edit vlans on page 63 *vlan-name*]

Release Information Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.

Description Specify a service identifier for each multi-chassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).

Options **number**—A number that identifies a particular service.

Range: 1 through 65535

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

static-mac

Syntax

```
static-mac mac-address {
  vlan-id number;
}
```

Hierarchy Level

[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name*bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name*],
[edit vlans on page 63 *vlan-name* switch-options interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.
[edit vlans *vlan-name* switch-options interface *interface name*] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers and EX Series switches only) Configure a static MAC address for a logical interface in a bridge domain or VLAN.

The **vlan-id** option can be specified for **static-macs** only if **vlan-id all** is configured for the bridging domain or VLAN.

Options

mac-address—MAC address

vlan-id *number*—(Optional) VLAN identifier to associate with static MAC address.

Required Privilege Level

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

Related Documentation

- Layer 2 Learning and Forwarding for Bridge Domains Overview
- Layer 2 Learning and Forwarding for VLANs Overview on page 4

swap

Syntax swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

On MX Series routers, you can enter this statement for Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet using Gigabit Ethernet IQ interfaces.

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

Related Documentation
- Rewriting the VLAN Tag on Tagged Frames

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.

On MX Series routers, this command can be used 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.

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

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag on page 39

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.

On MX Series routers, you can use this statement for 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.

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

switch-options

Syntax

```
switch-options {
  interface interface-name {
    interface-mac-limit limit {
      packet-action drop;
    }
    no-mac-learning;
    static-mac static-mac-address {
      vlan-id number;
    }
  }
  interface-mac-limit limit {
    packet-action drop;
  }
  mac-statistics;
  mac-table-size limit {
    packet-action drop;
  }
  no-mac-learning;
}
```

Hierarchy Level

[edit vlans on page 63 *vlan--name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* vlans on page 63 *vlan-name*],
[edit routing-instances *routing-instance-name* vlans on page 63 *vlan-name*]

Release Information

Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.

Description

Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.

The remaining statements are explained separately.

Required Privilege Level

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

tag-protocol-id (TPID to Rewrite)

Syntax tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers, value the default TPID value (**0x8100**) is supported.

Default If the **tag-protocol-id** statement is not configured, the TPID value is 0x8100.

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

Related Documentation
- Configuring Inner and Outer TPIDs and VLAN IDs

traceoptions (MVRP)

Syntax

```
traceoptions {
  file name <size size> <files number>  <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols mvrp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type),
[edit protocols mvrp],
[edit routing-instances *routing-instance-name* protocols mvrp] (for virtual switch instance type)

Release Information

Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description

For Multiple VLAN Registration Protocol (MVRP), configure tracing options.

Default

Traceoptions is disabled.

Options

disable —(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the **file** statement, you must specify a filename. Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place MVRP tracing output in the file **/var/log/mvrp-log**.

files *number*—(Optional) Maximum number of trace files, in the range from 2 through 1000. The default is 1 trace file. 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.

flag *flag*—Specify which tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Enable all trace options flags.
- **error**—Trace all failure conditions.
- **events**—Trace process state change and cleanup events.
- **pdu**—Trace RAPS PDU reception and transmission.
- **socket**—Trace socket activity.
- **state-machine**—Trace information about the state machine.
- **timers**—Trace protocol timers.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0.** When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0** . This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. The file size range is from 10240 through 4294967295. The default file size is 1MB.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

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

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP
- Configuring Multiple VLAN Registration Protocol (MVRP) on page 33
- Verifying That MVRP Is Working Correctly on page 129
- Understanding Multiple VLAN Registration Protocol (MVRP) on page 7

vlan-id (Bridge Domain or VLAN)

Syntax vlan-id (all | none | *number*);

Hierarchy Level [edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* bridge-domains*bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains*bridge-domain-name*],
[edit routing-instances *routing-instance-name*bridge-domains *bridge-domain-name*]
[edit vlans *vlan-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Support for Layer 2 trunk ports added in Junos OS Release 9.2.
Support for SRX 3400, SRX 3600, SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify a VLAN identifier (VID) to include in the packets sent to and from the bridge domain or VLAN, or a VPLS routing instance.

NOTE: When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the none option are not permitted.

Options *number*—A valid VLAN identifier. If you configure multiple bridge domains or VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for bridge domains or VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.

NOTE: If you specify a VLAN identifier, you cannot also use the all option. They are mutually exclusive.

all—Specify that the bridge domain or VLAN spans all the VLAN identifiers configured on the member logical interfaces.

NOTE: You cannot specify the all option if you include a routing interface in the bridge domain or VLAN.

none—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.

NOTE: Multichassis link aggregation (MC-LAG) does not support the **none** option with the **vlan-id** statement with bridge domains or VLANs.

Required Privilege Level

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

Related Documentation

- Configuring a Bridge Domain
- Configuring a VLAN on page 18
- Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances
- Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 18
- Configuring Bridge Domains as Switches for Layer 2 Trunk Ports
- Configuring a Layer 2 Virtual Switch
- Configuring a Layer 2 Virtual Switch on page 35
- Example: Configuring E-LINE and E-LAN Services for a PBB Network
- Example: Configuring Interfaces and Routing Instances for a User Logical System
- bridge-domains
- vlans on page 63

vlan-id (VLAN ID to Rewrite)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the **vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **vlan-id** statement that you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

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

Related Documentation

- Rewriting the VLAN Tag on Tagged Frames
- Binding VLAN IDs to Logical Interfaces

vlan-id-list

Syntax vlan-id-list [*vlan-id-numbers*];

Hierarchy Level
[edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]
[edit interfaces *interface-name* unit 0],
[edit vlans *vlan-name*]

Release Information
Statement introduced in Junos OS Release 9.4.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.

Specify the **trunk** option in the **interface-mode** statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id-list** statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

Options *vlan-id-numbers*—Valid VLAN identifiers. You can combine individual numbers with range lists including a hyphen.

Range: 0 through 4095

NOTE: On EX Series switches, the range is 0 through 4094.

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

Related Documentation

- Configuring a Bridge Domain
- Configuring a VLAN on page 18
- Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances
- Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 18

vlan-rewrite

Syntax vlan-rewrite translate (200 500 | 201 501)

Hierarchy Level [edit interfaces *interface-name* unit *number* family bridge interface-mode trunk]
[edit interfaces *interface-name* unit *number* family ethernet-switching interface-mode trunk]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.

Options

translate 200 500—Translates incoming packets with VLAN 200 to 500.

translate 201 501—Translates incoming packets with VLAN 201 to 501.

translate 202 502—Translates incoming packets with VLAN 202 to 502.

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

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag on page 39

vlan-tags

Syntax vlan-tags outer *number* inner *number*;

Hierarchy Level
[edit bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*],
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name*]
[edit vlans on page 63 *vlan-name*]

Release Information
Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify dual VLAN identifier tags for a bridge domain, VLAN, or VPLS routing instance.

Options
outer *number*—A valid VLAN identifier.

inner *number*—A valid VLAN identifier.

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

Related Documentation
- Configuring a Bridge Domain
- Configuring a VLAN on page 18
- Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances
- Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 18
- Configuring a Layer 2 Virtual Switch.
- Configuring a Layer 2 Virtual Switch on page 35

PART 3

Administration

CHAPTER 6

Routine Monitoring

Verifying That MVRP Is Working Correctly

Purpose After configuring your MX Series router or EX Series switch to participate in Multiple VLAN Registration Protocol (MVRP), verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.

Action

1. Confirm that the router is declaring Virtual LANs (VLANs).

   Show that MVRP is enabled:

   ```
   user@host> show mvrp
   MVRP configuration for routing instance 'default-switch'
   MVRP dynamic VLAN creation : Enabled
   MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
   MVRP timers (ms)
     Interface       Join   Leave  LeaveAll
     ge-11/3/0        200     800     10000
   ```

   Show the MVRP applicant state:

   ```
   user@host> show mvrp applicant-state
   MVRP applicant state for routing instance 'default-switch'
    (VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,

    (AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,

    (AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
    (AP) Anxious passive, (QP) Quiet passive

   VLAN Id     Interface          State
       100     ge-11/3/0          Declaring (QA)
       200     ge-11/3/0          Declaring (QA)
       300     ge-11/3/0          Declaring (QA)
   ```

2. Confirm that VLANs are registered on interfaces.

   List VLANs in the registered state:

   ```
   user@host> show mvrp registration-state
   MVRP registration state for routing instance 'default-switch'

   VLAN Id   Interface   Registrar   Forced      Managed    STP
                         State       State       State      State
       100   ge-11/3/0   Registered  Registered  Normal     Forwarding
   ```

```
    200   ge-11/3/0   Registered  Registered  Normal     Forwarding
    300   ge-11/3/0   Empty       Empty       Normal     Forwarding
```

3. Display a list of VLANs created dynamically.

   List dynamic VLAN membership:

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
 (s) static vlan, (f) fixed registration

VLAN Id            Interfaces
    100            ge-3/3/0
                   ge-3/0/5
    200            ge-3/3/0
                   ge-3/0/5
```

Meaning

The output of **show mvrp applicant-state** shows that trunk interface **ge-11/3/0** is declaring (sending out) interest in the VLAN IDs **100**, **200**, and **300** and MVRP is operating properly.

The output of **show mvrp registrant-state** shows the registrar state for VLANs **100** and **200** as **Registered**, indicating that these VLANs are receiving traffic from a customer site. VLAN **300** is in an **Empty** state and is not receiving traffic from a customer site.

The output of the **show mvrp dynamic-vlan-membership** shows that VLANs **100** and **200** are created dynamically (here, on an MX Series router operating as an aggregation switch between MX Series routers operating as edge switches). VLANs created statically are marked with an **(s)** (which is not indicated in this output).

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)
- Controlling the Management State of a VLAN in MVRP Configurations (CLI Procedure)

CHAPTER 7

Operational Commands

show ethernet-switching flood

Syntax

show ethernet-switching flood
<brief | detail | extensive>
<event-queue>
<instance *instance-name*>
<logical-system *logical-system-name*>
<route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)>
<vlan-name *vlan-name*>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers and EX Series switches only) Display Ethernet-switching flooding information.

Options

none—Display all Ethernet-switching flooding information for all VLANs.

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

event-queue—(Optional) Display the queue of pending Ethernet-switching flood events.

instance *instance-name*—(Optional) Display Ethernet-switching flooding information for the specified routing instance.

logical-system *logical-system-name*—(Optional) Display Ethernet-switching flooding information for the specified logical system.

route (all-ce-flood | all ve-flood | alt-root-flood | bd-flood | mlp-flood | re-flood)—(Optional) Display the following:

- **all-ce-flood**—Display the route for flooding traffic to all customer edge routers or switches if **no-local-switching** is enabled.
- **all-ve-flood**—Display the route for flooding traffic to all VPLS edge routers or switches if **no-local-switching** is enabled.
- **alt-root-flood**—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface.
- **bd-flood**—Display the route for flooding traffic of a VLAN if **no-local-switching** is not enabled.
- **mlp-flood**—Display the route for flooding traffic to MAC learning chips.
- **re-flood**—Display the route for Routing Engine flooding to all interfaces.

vlan-name *vlan-name*—(Optional) Display Ethernet-switching flooding information for the specified VLAN.

Required Privilege Level

view

List of Sample Output

show ethernet-switching flood on page 134

Sample Output

show ethernet-switching flood

```
user@host> show ethernet-switching flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__     comp            12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__    comp            12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__     comp            12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__    comp            12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__     comp            12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__    comp            12881
```

show ethernet-switching flood brief

```
user@host> show ethernet-switching flood brief
Name                Active CEs         Active VEs
__juniper_private1__ 0                 0
default-switch      9                  0
```

show ethernet-switching flood detail

```
user@host> show ethernet-switching flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__     comp            12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__    comp            12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__     comp            12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__    comp            12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                NhType          NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__     comp            12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__    comp            12881
```

show ethernet-switching flood extensive

```
user@host> show ethernet-switching flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
```

```
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
  Flood route prefix: 0x3057b/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12866
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12860
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605

  Flood route prefix: 0x30004/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12863
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12860
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605
VLAN Name: VLAN102

  Flood route prefix: 0x3057c/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12875
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12869
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605

  Flood route prefix: 0x30005/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12872
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12869
        Composition: split-horizon
```

```
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605
VLAN Name: VLAN103

  Flood route prefix: 0x3057d/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12884
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12878
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605

  Flood route prefix: 0x30006/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12881
    Flooding to:
    Name             Type          NhType          Index
    __all_ces__      Group          comp            12878
        Composition: split-horizon
        Flooding to:
        Name             Type          NhType          Index
        ae20.0           CE             ucst            7605
VLAN Name: VLAN104
```

show ethernet-switching interface

Syntax

show ethernet-switching interface
<brief | detail | extensive>
<interface-name>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers and EX Series switches only) Display Layer 2 learning information for all the interfaces.

Options

none—Display Ethernet-switching information for all interfaces.

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

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

Required Privilege Level

view

List of Sample Output

show ethernet-switching interface ae10.0 on page 138

Output Fields

Table 7 on page 137 describes the output fields for the **show ethernet-switching interface** command. Output fields are listed in the approximate order in which they appear.

Table 7: show ethernet-switching interface Output Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
VLAN members	VLANs associated with this interface.
Tag	VLAN ID.
MAC limit	Number of MAC addresses that can be associated with the interface.
STP state	Spanning tree protocol (STP) state.
Logical interface flags	Status of Layer 2 learning properties for each interface: • **DL**—MAC learning is disabled. • **LH**—MAC interface limit has been reached.. • **AD**—Packets are dropped after the MAC interface limit is reached. • **DN**—The MAC interface is down.
Tagging	Tagging state of the VLAN.

Sample Output

show ethernet-switching interface ae10.0

```
user@host> show ethernet-switching interface ae10.0
Logical Interface flags (DL - disable learning, AD - packet action drop,
                         LH - MAC limit hit, DN - interface down )
Logical      Vlan       TAG   MAC      STP          Logical         Tagging
interface    members          limit    state        interface flags
ae10.0                        8192                                  tagged
             VLAN70..   701
                               1024     Forwarding
             VLAN70..   702
                               1024     Forwarding
             VLAN70..   703
                               1024     Forwarding
             VLAN70..   704
                               1024     Forwarding
             VLAN70..   705
                               1024     Forwarding
             VLAN70..   706
                               1024     Forwarding
             VLAN70..   707
                               1024     Forwarding
             VLAN70..   708
                               1024     Forwarding
             VLAN70..   709
                               1024     Forwarding
             VLAN71..   710
                               1024     Forwarding
             VLAN71..   711
                               1024     Forwarding
             VLAN71..   712
                               1024     Forwarding
             VLAN71..   713
                               1024     Forwarding
             VLAN71..   714
                               1024     Forwarding
             VLAN71..   715
```

[...output truncated...]

show ethernet-switching statistics

Syntax

show ethernet-switching statistics
<**instance** *instance-name*>
<**logical-system** *logical-system-name*>
<**vlan-name** *vlan-name*>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers and EX Series switches only) Display Ethernet-switching statistics.

Options

none—Display Ethernet-switching statistics for all VLANs in all routing instances.

instance ***instance-name***—(Optional) Display statistics for the specified routing instance.

logical-system ***logical-system-name***—(Optional) Display Ethernet-switching statistics information for the specified logical system.

vlan-name ***vlan-name***—(Optional) Display statistics for the specified VLAN.

Required Privilege Level

view

List of Sample Output

show ethernet-switching statistics on page 140

Sample Output

show ethernet-switching statistics

```
user@host> show ethernet-switching statistics
 Local interface: ae1.0, Index: 1035
   Broadcast packets:                   220
   Broadcast bytes  :                 13720
   Multicast packets:                   130
   Multicast bytes  :                 11700
   Flooded packets  :                     0
   Flooded bytes    :                     0
   Unicast packets  :                     0
   Unicast bytes    :                     0
   Current MAC count:                     0 (Limit 1024)
 Local interface: vt-3/3/10.1048576, Index: 1280
   Broadcast packets:                     0
   Broadcast bytes  :                     0
   Multicast packets:                     0
   Multicast bytes  :                     0
   Flooded packets  :                     2
   Flooded bytes    :                   128
   Unicast packets  :                   632
   Unicast bytes    :                 39184
   Current MAC count:                     2
 Local interface: ge-3/1/2.0, Index: 1258
   Broadcast packets:                   100
   Broadcast bytes  :                  6800
   Multicast packets:                   200
   Multicast bytes  :                 18000
   Flooded packets  :                     0
   Flooded bytes    :                     0
   Unicast packets  :                   632
   Unicast bytes    :                 39184
   Current MAC count:                     2 (Limit 1024)
 Local interface: ae3.0, Index: 1043
   Broadcast packets:                     0
   Broadcast bytes  :                     0
   Multicast packets:                     0
   Multicast bytes  :                     0
   Flooded packets  :                     0
   Flooded bytes    :                     0
   Unicast packets  :                     0
   Unicast bytes    :                     0
   Current MAC count:                     0 (Limit 1024)
 Local interface: ge-3/3/8.0, Index: 1276
   Broadcast packets:                     0
   Broadcast bytes  :                     0
   Multicast packets:                     0
   Multicast bytes  :                     0
   Flooded packets  :                     0
   Flooded bytes    :                     0
   Unicast packets  :                     0
   Unicast bytes    :                     0
   Current MAC count:                     0 (Limit 8192)
 Local interface: ae5.0, Index: 1045
   Broadcast packets:                     0
   Broadcast bytes  :                     0
   Multicast packets:                     0
   Multicast bytes  :                     0
   Flooded packets  :                     0
   Flooded bytes    :                     0
```

```
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 8192)
Local interface: ae4.0, Index: 1044
  Broadcast packets:                   200
  Broadcast bytes  :                 13600
  Multicast packets:                     0
  Multicast bytes  :                     0
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 8192)
Local interface: ae26.0, Index: 1042
  Broadcast packets:                     0
  Broadcast bytes  :                     0
  Multicast packets:                     0
  Multicast bytes  :                     0
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                     0
  Unicast bytes    :                     0
  Current MAC count:                     0 (Limit 8192)
Local interface: ae25.0, Index: 1041
  Broadcast packets:                   133
  Broadcast bytes  :                  7980
  Multicast packets:                369934
  Multicast bytes  :              59207572
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                  1433
  Unicast bytes    :                119930
  Current MAC count:                     3 (Limit 8192)
Local interface: ae23.0, Index: 1040
  Broadcast packets:                   226
  Broadcast bytes  :                 14464
  Multicast packets:                585668
  Multicast bytes  :             153464476
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                 26552
  Unicast bytes    :               1947627
  Current MAC count:                     7 (Limit 8192)
Local interface: ae20.0, Index: 1039
  Broadcast packets:                   115
  Broadcast bytes  :                  6900
  Multicast packets:                395113
  Multicast bytes  :              61622869
  Flooded packets  :                     0
  Flooded bytes    :                     0
  Unicast packets  :                  1419
  Unicast bytes    :                117924
  Current MAC count:                     4 (Limit 8192)
```

show ethernet-switching table

Syntax

show ethernet-switching table
<brief | count | detail | extensive>
<*address*>
<instance *instance-name*>
<interface *interface-name*>
<interface *interface-name*>
isid *isid*
<*address*>
<vlan-id (all-vlan | *vlan-id*)>
<vlan-name (all | *vlan-name*)>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers and EX Series switches only) Display Layer 2 MAC address information.

Options

none—Display all learned Layer 2 MAC address information.

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

address—(Optional) Display the specified learned Layer 2 MAC address information.

instance *instance-name*—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.

interface *interface-name*—(Optional) Display learned Layer 2 MAC addresses for the specified interface.

isid *isid*—(Optional) Display learned Layer 2 MAC addresses for the specified ISID.

logical-system *logical-system-name*—(Optional) Display Ethernet-switching statistics information for the specified logical system.

vlan-id (all-vlan | *vlan-id*)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

vlan-name (all | *vlan-name*)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

Additional Information

When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

Required Privilege Level

view

List of Sample Output

show ethernet-switching table on page 144
show ethernet-switching table brief on page 145
show ethernet-switching table count on page 146

show ethernet-switching table extensive on page 147

Output Fields Table 8 on page 143 describes the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 8: show ethernet-switching table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
VLAN name	Name of the VLAN.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address is configured. • **D**—Dynamic MAC address is configured. • **L**—Locally learned MAC address is configured. • **SE**—MAC accounting is enabled. • **NM**—Non-configured MAC. • **R**—Locally learned MAC address is configured.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or VLAN in which the MAC address was learned.
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning-tree-protocolepoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show ethernet-switching table

```
user@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1102            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
```

```
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1103            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1104            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1105            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1106            00:1f:12:32:f5:c1   D             -    ae0.0
[...output truncated...]
```

show ethernet-switching table brief

```
user@host> show ethernet-switching table brief
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -    ae20.0
```

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags               interface
    VLAN104             88:e0:f3:bb:07:f0   D             -   ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags               interface
    VLAN1101            00:1f:12:32:f5:c1   D             -   ae0.0
[...output truncated...]
```

show ethernet-switching table count

```
user@host> show ethernet-switching table count
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000

1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           102              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           103              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN104
ae20.0:104

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           104              1                 0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108
```

```
0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
ae0.0:1101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1101             1                0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
ae0.0:1102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1102             1                0
[...output truncated...]
```

show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 101
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 102
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 103
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 104
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1101
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
```

```
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1102
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1103
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1104
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008
```

show interfaces (10-Gigabit Ethernet)

Syntax

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

Release Information

Command introduced in Junos OS Release 8.0.

Description

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

Options

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

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

descriptions—(Optional) Display interface description strings.

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

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

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 164
show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 167
show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 169
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 172
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 172
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 174

Output Fields

See Table 9 on page 150 for the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Table 9: show interfaces Gigabit Ethernet Output Fields

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

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

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

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

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 9 on page 150.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

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

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**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 9: 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 9: 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 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 10 on page 164 • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

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

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

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

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

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

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

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

Field Name	Field Description	Level of Output

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

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

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

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

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

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

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

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

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

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

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. Table 10 on page 164 describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In Table 10 on page 164, the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit **.50** (VLAN 50).

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

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

Sample Output

show interfaces extensive

```
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
```

(10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```
  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 accounting overhead: 100
  Ingress accounting 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

```
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
```

(10-Gigabit Ethernet, WAN PHY Mode)

```
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

```
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
```

(10-Gigabit Ethernet, DWDM OTN PIC)

```
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

```
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,
```

Mode, Unidirectional Mode, Transmit-Only)

```
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 accounting overhead: 100
    Ingress accounting 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

```
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,
```

Mode, Unidirectional Mode, Receive-Only)

```
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 (Gigabit Ethernet)

Syntax

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

Release Information

Command introduced before Junos OS Release 7.4.

Description

(M Series, T Series, and MX Series routers and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.

Options

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

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

descriptions—(Optional) Display interface description strings.

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

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

statistics—(Optional) Display static interface statistics.

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Related Documentation

- Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

List of Sample Output

show interfaces (Gigabit Ethernet) on page 192
show interfaces (Gigabit Ethernet on MX Series Routers) on page 192
show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 193
show interfaces brief (Gigabit Ethernet) on page 194
show interfaces detail (Gigabit Ethernet) on page 194
show interfaces extensive (Gigabit Ethernet IQ2) on page 195
show interfaces (Gigabit Ethernet Unnumbered Interface) on page 198
show interfaces (ACI Interface Set Configured) on page 199

Output Fields

Table 11 on page 177 describes the output fields for the **show interfaces** (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 12 on page 190.

Table 11: 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 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 11: 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 11: 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 router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
Interface transmit statistics	(On MX Series devices) Status of the **interface-transmit-statistics** configuration: Enabled or Disabled. • **Enabled**—When the **interface-transmit-statistics** statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • **Disabled**—When the **interface-transmit-statistics** statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics: • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—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 the PCS receiver is operating in normal mode.	**detail extensive**

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—There are two possible conditions regarding the number of oversized frames: ◦ Packet length exceeds 1518 octets, or ◦ Packet length exceeds MRU • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the **VLAN tagged frames** field displays 0 when the **show interfaces** command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

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

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

Table 11: 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 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	**brief detail extensive** none
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under Common Output Fields Description.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	**detail extensive** none
Service Name Table	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	**detail extensive** none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	**detail extensive** none

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: **On** or **Off**. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the router.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under Common Output Fields Description.	**detail extensive**
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	**detail extensive** none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	**detail extensive**

Table 11: show interfaces Gigabit Ethernet Output Fields *(continued)*

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

Table 12: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

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

Table 12: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type *(continued)*

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Outbound logical interface	**show interfaces ge-0/0/0.50 extensive**	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
 Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
    Flags: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Input packets : 0
    Output packets: 0
     Protocol aenet, AE bundle: ae0.0    Link Index: 4
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
   Output bytes  :         240614363944            772721536 bps
   Output packets:           3538446506              1420444 pps
    Direction : Output
```

statistics configuration)

```
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
     Output bytes  :         195560312716            522726272 bps
     Output packets:           4251311146              1420451 pps
```

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
    0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Drop   bytes  :                    0                    0 bps
   Drop   packets:                    0                    0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0
```

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

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
      Flags: Is-Primary

  Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
  (Generation 139)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
```

show interfaces extensive

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
```

(Gigabit Ethernet IQ2)

```
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags     : None
CoS queues     : 8 supported, 4 maximum usable queues
Schedulers     : 256
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped   : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :          38910844056                 7952 bps
 Output bytes  :              7174605                 8464 bps
 Input  packets:            418398473                   11 pps
 Output packets:                78903                   12 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Ingress traffic statistics at Packet Forwarding Engine:
 Input  bytes  :          38910799145                 7952 bps
 Input  packets:            418397956                   11 pps
 Drop   bytes  :                   0                    0 bps
 Drop   packets:                   0                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort            418390823            418390823                    0

  1 expedited-fo                 0                    0                    0

  2 assured-forw                 0                    0                    0

  3 network-cont              7133                 7133                    0

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

  0 best-effort               1031                 1031                    0

  1 expedited-fo                 0                    0                    0

  2 assured-forw                 0                    0                    0

  3 network-cont             77872                77872                    0

Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                  38910844056          7174605
```

```
    Total packets                    418398473            78903
    Unicast packets                408021893366             1026
    Broadcast packets                       10               12
    Multicast packets                418398217            77865
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0  OTN Received Overhead Bytes:
    APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
    Payload Type: 0x08
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x08
  Filter statistics:
    Input packet count               418398473
    Input packet rejects                   479
    Input DA rejects                       479
    Input SA rejects                         0
    Output packet count                                   78903
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
        Remote fault: OK
    Local resolution:
        Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                                %            bps     %           usec
    0 best-effort              95      950000000    95              0
low    none
    3 network-control           5       50000000     5              0
low    none
    Direction : Input
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                                %            bps     %           usec
    0 best-effort              95      950000000    95              0
low    none
    3 network-control           5       50000000     5              0
low    none

  Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               812400
     Output bytes  :              1349206
     Input  packets:                 9429
     Output packets:                 9449
     IPv6 transit statistics:
      Input  bytes  :                   0
```

```
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Local statistics:
   Input  bytes  :                812400
   Output bytes  :               1349206
   Input  packets:                  9429
   Output packets:                  9449
  Transit statistics:
   Input  bytes  :                   0                 7440 bps
   Output bytes  :                   0                 7888 bps
   Input  packets:                   0                   10 pps
   Output packets:                   0                   11 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Protocol inet, MTU: 1500, Generation: 169, Route table: 0
    Flags: Is-Primary, Mac-Validate-Strict
    Mac-Validate Failures: Packets: 0, Bytes: 0
    Addresses, Flags: Is-Preferred Is-Primary
    Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
    Output Filters: F2-ge-3/0/1.0-out (53)
    Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
      Generation: 196
  Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
    Flags: Is-Primary
    Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in Table 11 on page 177.

show interfaces (Gigabit Ethernet

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
```


Unnumbered Interface)

```
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues     : 8 supported, 4 maximum usable queues
Current address: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 624 bps (1 pps)
Active alarms  : None
Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
    Flags: Unnumbered
    Donor interface: lo0.0 (Index 64)
    Preferred source address: 22.22.22.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
  Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ]  Encapsulation: PPP-over-

    Ethernet
    ACI VLAN:
      Dynamic Profile: aci-vlan-set-profile
    PPPoE:
      Dynamic Profile: aci-vlan-pppoe-profile,
      Service Name Table: None,
      Max Sessions: 32000, Max Sessions VSA Ignore: Off,
      Duplicate Protection: On, Short Cycle Protection: Off,
      AC Name: nbc
    Input packets : 9
    Output packets: 8
    Protocol multiservice, MTU: Unlimited
```

show interfaces irb

Syntax

show interfaces irb
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance *instance-name*>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

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.

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

routing-instance *instance-name*—(Optional) Display information for the interface with the specified SNMP index.

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 VLAN that has a Layer 3 protocol configured.

Required Privilege Level

view

List of Sample Output

show interfaces irb extensive on page 205
show interfaces irb snmp-index on page 206

Output Fields

Table 13 on page 200 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 13: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the "Enabled Field" section under Common Output Fields Description.	All levels

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

Field Name	Field Description	Level of Output
Proto	Protocol configured on the interface.	**terse**
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Physical interface type.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**detail extensive brief** none
MTU	MTU size on the physical interface.	**detail extensive brief** none
Clocking	Reference clock source: **Internal** or **External**. Always unspecified on IRB interfaces.	**detail extensive brief**
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	**detail extensive brief**
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under Common Output Fields Description.	**detail extensive brief** none
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under Common Output Fields Description.	**detail extensive brief** none
Link type	Physical interface link type: **full duplex** or **half duplex**.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Links Flags" section under Common Output Fields Description.	**detail extensive** none
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Alternate link address	Backup address of the link.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 13: 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 13: 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 13: 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 irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
    Protocol multiservice, MTU: 1500
      Flags: Is-Primary
```

show interfaces queue

Syntax

show interfaces queue
<aggregate | remaining-traffic>
<both-ingress-egress>
<egress>
<forwarding-class *forwarding-class*>
<ingress>
<interface-name *interface-name*>
<l2-statistics>
<remaining-traffic>

Release Information

Command introduced before Junos OS Release 7.4.
both-ingress-egress, **egress**, and **ingress** options introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 11.1 for the QFX Series.
l2-statistics option introduced in Junos OS Release 12.1.

Description

Display class-of-service (CoS) queue information for physical interfaces.

Options

none—Show detailed CoS queue statistics for all physical interfaces.

aggregate—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

egress—(Optional) Display egress queue statistics.

forwarding-class *forwarding-class*—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

interface-name *interface-name*—(Optional) Show detailed CoS queue statistics for the specified interface.

l2-statistics—(optional) Display layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in Table 14 on page 208. Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the layer 3 level. In the case of link fragmentation and interleaving (LFI) for which not fragmentation is applied, corresponding layer 2 overheads are added, as shown in Table 14 on page 208.

Table 14: Layer 2 Overhead, Transmitted Packets/Bytes

Protocol	Fragmentation		LFI
	First fragmentation	Second to n fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Layer 2 Statistics - Fragmentation Overhead Calculation

```
MLPPP/MC-MLPPP Overhead details:
   ===============================
   Fragment 1:

    Outer PPP header                          : 4 bytes
    Long or short sequence MLPPP header       : 4 bytes or 2 bytes
    Inner PPP header                          : 1 byte
    HDLC flag and FCS bytes                   : 4 bytes

   Fragments 2 .. n :

    Outer PPP header                          : 4 bytes
    Long or short sequence MLPPP header       : 4 bytes or 2 bytes
    HDLC flag and FCS bytes                   : 4 bytes


   MLFR (FRF15) Overhead details:
   ============================
   Fragment 1:
       Framerelay header     :  2 bytes
       Control,NLPID         :  2 bytes
       Fragmentaion header   :  2 bytes
       Inner proto           :  2 bytes
       HDLC flag and  FCS    :  4 bytes

   Fragments 2 ...n :

       Framerelay header     :  2 bytes
       Control,NLPID         :  2 bytes
       Fragmentaion header   :  2 bytes
       HDLC flag and  FCS    :  4 bytes


   MFR (FRF16) Overhead details:
   ============================
   Fragment 1:
      Fragmentaion header   :  2 bytes
      Framerelay header     :  2 bytes
      Inner proto           :  2 bytes
      HDLC flag and  FCS    :  4 bytes

   Fragments 2 ...n :
      Fragmentaion header   :  2 bytes
      Framerelay header     :  2 bytes
      HDLC flag and  FCS    :  4 bytes
```

Overhead with LFI

```
MLPPP(Long & short sequence):
   ===========================
       Outer PPP header      :  4 bytes
       HDLC flag and  FCS    :  4 bytes

  MLFR (FRF15):
  =============
       Framerelay header     :  2 bytes
       Control,NLPID         :  2 bytes
       HDLC flag and  FCS    :  4 bytes
```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the layer 2 level, bytes transmitted is 1008 in 1 packet.

remaining-traffic—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

Additional Information

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the Junos® OS Network Interfaces. For related CoS operational mode commands, see the Junos OS Operational Mode Commands.

Required Privilege Level

view

List of Sample Output

show interfaces queue (Aggregated Ethernet on a T320 Router) on page 216
show interfaces queue (Fast Ethernet on a J4300 Router) on page 217

show interfaces queue (Gigabit Ethernet on a T640 Router) on page 218
show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC) on page 218
show interfaces queue (Gigabit Ethernet IQ2 PIC) on page 222
show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC) on page 225
show interfaces queue ingress (Gigabit Ethernet IQ2 PIC) on page 228
show interfaces queue egress (Gigabit Ethernet IQ2 PIC) on page 228
show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC) on page 230
show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode) on page 233
show interfaces queue (QFX Series) on page 243
show interfaces queue l2-statistics (lsq interface) on page 244

Output Fields Table 15 on page 211 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 15: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under Common Output Fields Description.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Ingress queues supported	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
Ingress queues in use	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
Output queues supported	Total number of output queues supported on the specified interface.
Output queues in use	Total number of output queues in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
Queue	Queue number.

Table 15: show interfaces queue Output Fields *(continued)*

Field Name	Field Description
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.
Burst size	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue. NOTE: For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type. For more information, see Table 16 on page 214.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values. NOTE: For layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 207
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by PIC type. For more information, see Table 16 on page 214. NOTE: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation. NOTE: For layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 207
Tail-dropped packets	Number of packets dropped because of tail drop.

Table 15: show interfaces queue Output Fields *(continued)*

Field Name	Field Description

RED-dropped packets	Number of packets dropped because of random early detection (RED).
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories:
	  • **Low, non-TCP**—Number of low-loss priority non-TCP packets dropped because of RED.
	  • **Low, TCP**—Number of low-loss priority TCP packets dropped because of RED.
	  • **High, non-TCP**—Number of high-loss priority non-TCP packets dropped because of RED.
	  • **High, TCP**—Number of high-loss priority TCP packets dropped because of RED.
	• (J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories:
	  • **Low**—Number of low-loss priority packets dropped because of RED.
	  • **Medium-low**—Number of medium-low loss priority packets dropped because of RED.
	  • **Medium-high**—Number of medium-high loss priority packets dropped because of RED.
	  • **High**—Number of high-loss priority packets dropped because of RED.

RED-dropped bytes	Number of bytes dropped because of RED. The byte counts vary by PIC type. For more information, see Table 16 on page 214.
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories:
	  • **Low, non-TCP**—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • **Low, TCP**—Number of low-loss priority TCP bytes dropped because of RED.
	  • **High, non-TCP**—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • **High, TCP**—Number of high-loss priority TCP bytes dropped because of RED.
	• (J Series routers only) The output classifies dropped bytes into the following categories:
	  • **Low**—Number of low-loss priority bytes dropped because of RED.
	  • **Medium-low**—Number of medium-low loss priority bytes dropped because of RED.
	  • **Medium-high**—Number of medium-high loss priority bytes dropped because of RED.
	  • **High**—Number of high-loss priority bytes dropped because of RED.

	Byte counts vary by PIC type. Table 16 on page 214 shows how the byte counts on the outbound interfaces vary depending on the PIC type. Table 16 on page 214 is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 16: Byte Count by PIC Type

PIC Type	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes	The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type. For RED dropped, 6 bytes are added for the source MAC address.
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet	–
Non-IQ PIC	Interface	T Series, TX Series, T1600, and MX Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. T4000 routers with Type 5 FPCs : • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. M Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. PTX Series Packet Transport Switches: • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). • RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted.	The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.

Table 16: Byte Count by PIC Type *(continued)*

PIC Type	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	T Series, TX Series, T1600, and MX Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. M Series routers: • Queued: 478 bytes of Layer 3 packet. • Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes • RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2–E PICs	Interface	Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs 8-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet.	–

Sample Output

show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets              :                    5                    0 pps
    Bytes                :                  242                    0 bps
  Transmitted:
    Packets              :                    5                    0 pps
    Bytes                :                  242                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets              :             42603765               595484 pps
    Bytes                :           5453281920            609776496 bps
  Transmitted:
    Packets              :             42603765               595484 pps
    Bytes                :           5453281920            609776496 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef1
  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: 3, Forwarding classes: nc
  Queued:
    Packets              :                   45                    0 pps
    Bytes                :                 3930                    0 bps
  Transmitted:
    Packets              :                   45                    0 pps
    Bytes                :                 3930                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 4, Forwarding classes: af11
  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: 5, Forwarding classes: ef11
```

```
  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: 6, Forwarding classes: af12
  Queued:
    Packets              :             31296413               437436 pps
    Bytes                :           4005940864            447935200 bps
  Transmitted:
    Packets              :             31296413               437436 pps
    Bytes                :           4005940864            447935200 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 7, Forwarding classes: nc2
  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 queue (Fast Ethernet on a J4300 Router)

```
user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets              :              5240762                 3404 pps
    Bytes                :           3020710354             15934544 bps
  Transmitted:
    Packets              :              5240762                 3404 pps
    Bytes                :           3020710354             15934544 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 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets              :              2480391                 1650 pps
    Bytes                :           1304685666              6945704 bps
  Transmitted:
    Packets              :              2478740                 1650 pps
    Bytes                :           1303817240              6945704 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                 1651                    0 pps
```

```
    Low                 :                    0                    0 pps
    Medium-low          :                    0                    0 pps
    Medium-high         :                    0                    0 pps
    High                :                 1651                    0 pps
   RED-dropped bytes    :               868426                    0 bps
    Low                 :                    0                    0 pps
    Medium-low          :                    0                    0 pps
    Medium-high         :                    0                    0 pps
    High                :               868426                    0 pps
```

show interfaces queue (Gigabit Ethernet on a T640 Router)

```
user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
   Interface index: 150, SNMP ifIndex: 42
 Forwarding classes: 8 supported, 8 in use
 Output queues: 8 supported, 8 in use
 Queue: 0, Forwarding classes: be
   Queued:
     Packets              :                   13                    0 pps
     Bytes                :                  622                    0 bps
   Transmitted:
     Packets              :                   13                    0 pps
     Bytes                :                  622                    0 bps
     Tail-dropped packets :                    0                    0 pps
     RED-dropped packets  :                    0                    0 pps
     RED-dropped bytes    :                    0                    0 bps
 Queue: 1, Forwarding classes: af1
   Queued:
     Packets              :           1725947945               372178 pps
     Bytes                :         220921336960            381110432 bps
   Transmitted:
     Packets              :           1725947945               372178 pps
     Bytes                :         220921336960            381110432 bps
     Tail-dropped packets :                    0                    0 pps
     RED-dropped packets  :                    0                    0 pps
     RED-dropped bytes    :                    0                    0 bps
 Queue: 2, Forwarding classes: ef1
   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: 3, Forwarding classes: nc
   Queued:
     Packets              :                  571                    0 pps
     Bytes                :                49318                  336 bps
   Transmitted:
     Packets              :                  571                    0 pps
     Bytes                :                49318                  336 bps
     Tail-dropped packets :                    0                    0 pps
     RED-dropped packets  :                    0                    0 pps
     RED-dropped bytes    :                    0                    0 bps
```

show interfaces queue aggregate (Gigabit

```
user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
```

Ethernet Enhanced DPC)

```
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             148450735               947295 pps
    Bytes                :            8016344944            409228848 bps
  Transmitted:
    Packets              :              76397439               487512 pps
    Bytes                :            4125461868            210602376 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :              72053285               459783 pps
     Low                 :              72053285               459783 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :            3890877444            198626472 bps
     Low                 :            3890877444            198626472 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 : Not Available
    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              :             410278257               473940 pps
    Bytes                :           22156199518            204742296 bps
  Transmitted:
    Packets              :               4850003                 4033 pps
    Bytes                :             261900162              1742256 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :             405425693               469907 pps
     Low                 :             405425693               469907 pps
     Medium-low          :                     0                    0 pps
     Medium-high         :                     0                    0 pps
     High                :                     0                    0 pps
    RED-dropped bytes    :           21892988124            203000040 bps
     Low                 :           21892988124            203000040 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 : Not Available
    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
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              76605230                 485376 pps
    Bytes                :            5209211400              264044560 bps
  Transmitted:
    Packets              :              76444631                 484336 pps
    Bytes                :            5198235612              263478800 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                160475                   1040 pps
     Low                 :                160475                   1040 pps
     Medium-low          :                     0                      0 pps
     Medium-high         :                     0                      0 pps
     High                :                     0                      0 pps
    RED-dropped bytes    :              10912300                 565760 bps
     Low                 :              10912300                 565760 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 : Not Available
    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              :               4836136                   3912 pps
    Bytes                :             333402032                2139056 bps
  Transmitted:
    Packets              :               3600866                   1459 pps
    Bytes                :             244858888                 793696 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :               1225034                   2450 pps
     Low                 :               1225034                   2450 pps
```

```
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :                    0                    0 pps
    RED-dropped bytes    :             83302312              1333072 bps
     Low                 :             83302312              1333072 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 : Not Available
    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

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             77059796               486384 pps
    Bytes                :           3544750624            178989576 bps
  Transmitted:
    Packets              :             77059797               486381 pps
    Bytes                :           3544750670            178988248 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: 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              :              4846580                 3934 pps
    Bytes                :            222942680              1447768 bps
  Transmitted:
    Packets              :              4846580                 3934 pps
    Bytes                :            222942680              1447768 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: 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 queue (Gigabit Ethernet IQ2 PIC)

```
user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            418390039                   10 pps
    Bytes                :          38910269752                 7440 bps
  Transmitted:
    Packets              :            418390039                   10 pps
    Bytes                :          38910269752                 7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    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 : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets             :                     0                     0 pps
    Bytes               :                     0                     0 bps
  Transmitted:
    Packets             :                     0                     0 pps
    Bytes               :                     0                     0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets             :                  7055                     1 pps
    Bytes               :                451552                   512 bps
  Transmitted:
    Packets             :                  7055                     1 pps
    Bytes               :                451552                   512 bps
    Tail-dropped packets : Not Available
    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             :                  1031                     0 pps
    Bytes               :                143292                     0 bps
  Transmitted:
    Packets             :                  1031                     0 pps
    Bytes               :                143292                     0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     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 : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets             :                     0                     0 pps
    Bytes               :                     0                     0 bps
  Transmitted:
    Packets             :                     0                     0 pps
    Bytes               :                     0                     0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                     0                     0 pps
    RL-dropped bytes     :                     0                     0 bps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
```

```
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                77009                   11 pps
    Bytes                :              6894286                 7888 bps
  Transmitted:
    Packets              :                77009                   11 pps
    Bytes                :              6894286                 7888 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                 1031                    0 pps
    Bytes                :               147328                    0 bps
  Transmitted:
    Packets              :                 1031                    0 pps
    Bytes                :               147328                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low, non-TCP        :                    0                    0 pps
     Low, TCP            :                    0                    0 pps
     High, non-TCP       :                    0                    0 pps
     High, TCP           :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low, non-TCP        :                    0                    0 bps
     Low, TCP            :                    0                    0 bps
     High, non-TCP       :                    0                    0 bps
     High, TCP           :                    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, non-TCP        :                    0                    0 pps
     Low, TCP            :                    0                    0 pps
     High, non-TCP       :                    0                    0 pps
     High, TCP           :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low, non-TCP        :                    0                    0 bps
     Low, TCP            :                    0                    0 bps
     High, non-TCP       :                    0                    0 bps
     High, TCP           :                    0                    0 bps
Queue: 2, Forwarding classes: assured-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, non-TCP        :                    0                    0 pps
```

```
      Low, TCP              :                       0                       0 pps
      High, non-TCP         :                       0                       0 pps
      High, TCP             :                       0                       0 pps
    RED-dropped bytes       :                       0                       0 bps
      Low, non-TCP          :                       0                       0 bps
      Low, TCP              :                       0                       0 bps
      High, non-TCP         :                       0                       0 bps
      High, TCP             :                       0                       0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                 :                   94386                      12 pps
    Bytes                   :                13756799                    9568 bps
  Transmitted:
    Packets                 :                   94386                      12 pps
    Bytes                   :                13756799                    9568 bps
    Tail-dropped packets    :                       0                       0 pps
    RED-dropped packets     :                       0                       0 pps
      Low, non-TCP          :                       0                       0 pps
      Low, TCP              :                       0                       0 pps
      High, non-TCP         :                       0                       0 pps
      High, TCP             :                       0                       0 pps
    RED-dropped bytes       :                       0                       0 bps
      Low, non-TCP          :                       0                       0 bps
      Low, TCP              :                       0                       0 bps
      High, non-TCP         :                       0                       0 bps
      High, TCP             :                       0                       0 bps
```

show interfaces queue both-ingress-egress

```
user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
```

(Gigabit Ethernet IQ2 PIC)

```
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                  254                    0 pps
    Bytes                :                16274                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    3                    0 pps
    Bytes                :                  126                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
```

```
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                     0                    0 bps
  Transmitted:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                     0                    0 bps
  Transmitted:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              80564692                    0 pps
    Bytes                :            3383717100                    0 bps
  Transmitted:
    Packets              :              80564692                    0 pps
    Bytes                :            3383717100                    0 bps
    Tail-dropped packets :                     0                    0 pps
    RED-dropped packets  :                     0                    0 pps
    RED-dropped bytes    :                     0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :              80564685                    0 pps
    Bytes                :            3383716770                    0 bps
  Transmitted:
    Packets              :              80564685                    0 pps
    Bytes                :            3383716770                    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              :                     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: 3, Forwarding classes: network-control
  Queued:
    Packets              :                  9397                    0 pps
    Bytes                :               3809052                  232 bps
  Transmitted:
    Packets              :                  9397                    0 pps
```

```
    Bytes                :              3809052                  232 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```
user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                  288                    0 pps
    Bytes                :                18450                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
```

show interfaces queue egress

```
user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
```

(Gigabit Ethernet IQ2 PIC)

```
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    3                    0 pps
    Bytes                :                  126                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             80564692                    0 pps
    Bytes                :           3383717100                    0 bps
  Transmitted:
    Packets              :             80564692                    0 pps
    Bytes                :           3383717100                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :             80564685                    0 pps
    Bytes                :           3383716770                    0 bps
  Transmitted:
    Packets              :             80564685                    0 pps
    Bytes                :           3383716770                    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              :                     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: 3, Forwarding classes: network-control
  Queued:
    Packets              :                  9538                     0 pps
    Bytes                :               3819840                     0 bps
  Transmitted:
    Packets              :                  9538                     0 pps
    Bytes                :               3819840                     0 bps
    Tail-dropped packets :                     0                     0 pps
    RED-dropped packets  :                     0                     0 pps
    RED-dropped bytes    :                     0                     0 bps
```

show interfaces queue remaining-traffic

```
user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
```

(Gigabit Ethernet Enhanced DPC)

```
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :             110208969               472875 pps
    Bytes                :            5951284434            204282000 bps
  Transmitted:
    Packets              :             110208969               472875 pps
    Bytes                :            5951284434            204282000 bps
    Tail-dropped packets : Not Available
    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: 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 : Not Available
    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              :                     0                    0 pps
    Bytes                :                     0                    0 bps
  Transmitted:
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets : Not Available
    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: 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 : Not Available
    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
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              109355853                 471736 pps
    Bytes                :             7436199152              256627968 bps
  Transmitted:
    Packets              :              109355852                 471736 pps
    Bytes                :             7436198640              256627968 bps
    Tail-dropped packets : Not Available
    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: 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 : Not Available
    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              :                      0                      0 pps
    Bytes                :                      0                      0 bps
  Transmitted:
    Packets              :                      0                      0 pps
    Bytes                :                      0                      0 bps
    Tail-dropped packets : Not Available
    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: 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  : Not Available
    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 queue (Channelized OC12 IQE

```
user@host> show interfaces queue t3-1/1/0:7
 Physical interface: t3-1/1/0:7, Enabled, Physical link is Up
```

Type 3 PIC in SONET Mode)

```
  Interface index: 192, SNMP ifIndex: 1948

  Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

 Forwarding classes: 16 supported, 9 in use

 Egress queues: 8 supported, 8 in use

 Queue: 0, Forwarding classes: DEFAULT

   Queued:

     Packets              :                214886                 13449 pps

     Bytes                :               9884756               5164536 bps

   Transmitted:

     Packets              :                214886                 13449 pps

     Bytes                :               9884756               5164536 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: 1, Forwarding classes: REALTIME

   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: PRIVATE
  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: 3, Forwarding classes: CONTROL
  Queued:
    Packets              :                60                 0 pps
```

```
    Bytes                :                 4560                 0 bps
  Transmitted:
    Packets              :                   60                 0 pps
    Bytes                :                 4560                 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: 4, Forwarding classes: CLASS_B_OUTPUT
  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: 5, Forwarding classes: CLASS_C_OUTPUT
  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: 6, Forwarding classes: CLASS_V_OUTPUT
  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: 7, Forwarding classes: CLASS_S_OUTPUT, GETS
  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


Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT
```

```
  Queued:
    Packets              :              371365              23620 pps
    Bytes                :            15597330            7936368 bps
  Transmitted:
    Packets              :              371365              23620 pps
    Bytes                :            15597330            7936368 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: 1, Forwarding classes: REALTIME
  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: PRIVATE
  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: 3, Forwarding classes: CONTROL
  Queued:
    Packets             :                32843                    0 pps
    Bytes               :              2641754                   56 bps
  Transmitted:
    Packets             :                32843                    0 pps
    Bytes               :              2641754                   56 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: 4, Forwarding classes: CLASS_B_OUTPUT
  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: 5, Forwarding classes: CLASS_C_OUTPUT
  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: 6, Forwarding classes: CLASS_V_OUTPUT
  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: 7, Forwarding classes: CLASS_S_OUTPUT, GETS
  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 queue (QFX Series)

```
user@switch> show interfaces queue xe-0/0/15
   Physical interface: xe-0/0/15, Enabled, Physical link is Up
     Interface index: 49165, SNMP ifIndex: 539
   Forwarding classes: 12 supported, 8 in use
   Egress queues: 12 supported, 8 in use
   Queue: 0, Forwarding classes: best-effort
     Queued:
       Packets              :                    0                    0 pps
       Bytes                :                    0                    0 bps
     Transmitted:
       Packets              :                    0                    0 pps
       Bytes                :                    0                    0 bps
       Tail-dropped packets : Not Available
       Total-dropped packets:                    0                    0 pps
       Total-dropped bytes  :                    0                    0 bps
   Queue: 3, Forwarding classes: fcoe
     Queued:
       Packets              :                    0                    0 pps
```

```
      Bytes               :                    0                    0 bps
    Transmitted:
      Packets             :                    0                    0 pps
      Bytes               :                    0                    0 bps
      Tail-dropped packets : Not Available
      Total-dropped packets:                   0                    0 pps
      Total-dropped bytes  :                   0                    0 bps
0 bps
  Queue: 4, Forwarding classes: no-loss
    Queued:
      Packets             :                    0                    0 pps
      Bytes               :                    0                    0 bps
    Transmitted:
      Packets             :                    0                    0 pps
      Bytes               :                    0                    0 bps
      Tail-dropped packets : Not Available
      Total-dropped packets:                   0                    0 pps
      Total-dropped bytes  :                   0                    0 bps
  Queue: 7, 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 : Not Available
      Total-dropped packets:                   0                    0 pps
      Total-dropped bytes  :                   0                    0 bps
  Queue: 8, Forwarding classes: mcast
    Queued:
      Packets             :                    0                    0 pps
      Bytes               :                    0                    0 bps
    Transmitted:
      Packets             :                    0                    0 pps
      Bytes               :                    0                    0 bps
      Tail-dropped packets : Not Available
      Total-dropped packets:                   0                    0 pps
      Total-dropped bytes  :                   0                    0 bps
```

show interfaces queue l2-statistics (lsq interface)

```
user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets             :                    1                    0 pps
    Bytes               :                 1001                    0 bps
  Transmitted:
    Packets             :                    5                     0 pps
    Bytes               :                 1062                    0 bps
    Tail-dropped packets :                   0                    0 pps
    RED-dropped packets  :                   0                    0 pps
    RED-dropped bytes    :                   0                    0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets             :                    1                    0 pps
    Bytes               :                 1500                    0 bps
  Transmitted:
    Packets             :                    6                     0 pps
    Bytes               :                 1573                    0 bps
```

```
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets              :                    1                    0 pps
    Bytes                :                  512                    0 bps
  Transmitted:
    Packets              :                    3                    0 pps
    Bytes                :                  549                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: nc
  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 mvrp

Syntax show mvrp

Release Information Command introduced in Junos OS Release 10.1.

Description For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) configuration information.

Required Privilege Level view

Related Documentation

- show mvrp applicant-state on page 248
- show mvrp dynamic-vlan-memberships on page 250
- show mvrp interface on page 251
- show mvrp registration-state on page 252
- show mvrp statistics on page 254

List of Sample Output show mvrp on page 247

Output Fields Table 17 on page 246 lists the output fields for the **show mvrp** command. Output fields are listed in the approximate order in which they appear.

Table 17: show mvrp Output Fields

Field Name	Field Description
MVRP dynamic VLAN creation	Displays whether global MVRP dynamic Virtual LAN (VLAN) creation is **Enabled** or **Disabled**.
MVRP BPDU MAC address	Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.
MVRP timers (ms)	Displays MVRP timer information: • **Interface**—The interface on which MVRP is configured. • **Join**—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • **Leave**—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • **LeaveAll**—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.

Sample Output

show mvrp

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface       Join   Leave  LeaveAll
  ge-11/2/8        200     800     10000
  ge-11/0/9        200     800     10000
  ge-11/3/0        200     800     10000
```

show mvrp applicant-state

Syntax show mvrp applicant-state

Release Information Command introduced in Junos OS Release 10.1.

Description For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) applicant state information.

Required Privilege Level view

Related Documentation

- show mvrp on page 246
- show mvrp interface on page 251
- show mvrp registration-state on page 252
- show mvrp statistics on page 254

List of Sample Output show mvrp applicant-state on page 249

Output Fields Table 18 on page 248 lists the output fields for the **show mvrp applicant-state** command. Output fields are listed in the approximate order in which they appear.

Table 18: show mvrp applicant-state Output Fields

Field Name	Field Description
VLAN Id	Displays the Virtual LAN (VLAN) ID number.
Interface	Displays the interface number associated with the VLAN ID.
State	Displays one of the following MVRP registrar states: • **VO**—Very anxious observer. • **VP**—Very anxious passive. • **VA**—Very anxious new. • **AN**—Anxious new. • **AA**—Anxious active. • **QA**—Quiet active. • **LA**—Leaving active. • **AO**—Anxious observer. • **QO**—Quiet observer. • **LO**—Leaving observer. • **AP**—Anxious passive. • **QA**—Quiet passive.

Sample Output

show mvrp applicant-state

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
 (VO) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
 (AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
 (AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
 (AP) Anxious passive, (QP) Quiet passive

VLAN Id     Interface          State
    100     ge-11/3/0          Declaring (QA)
    200     ge-11/3/0          Declaring (QA)
    300     ge-11/3/0          Declaring (QA)
```

show mvrp dynamic-vlan-memberships

Syntax show mvrp dynamic-vlan-memberships

Release Information Command introduced in Junos OS Release 10.1.

Description For MX Series routers and EX Series switches, display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router or switch.

Required Privilege Level clear

Related Documentation

- show mvrp on page 246
- show mvrp applicant-state on page 248
- show mvrp interface on page 251
- show mvrp registration-state on page 252
- show mvrp statistics on page 254

List of Sample Output show mvrp dynamic-vlan-memberships on page 250

Output Fields Table 19 on page 250 lists the output fields for the **show mvrp dynamic-vlan-memberships** command. Output fields are listed in the approximate order in which they appear.

Table 19: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Id	The VLAN ID of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

Sample Output

show mvrp dynamic-vlan-memberships

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
 (s) static vlan, (f) fixed registration

VLAN Id            Interfaces
    100 (s)        ge-11/3/0
    200 (s)        ge-11/3/0
    300 (s)
```

show mvrp interface

Syntax show mvrp interface

Release Information Command introduced in Junos OS Release 10.1.

Description For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) interface-specific information.

Required Privilege Level view

Related Documentation

- show mvrp on page 246
- show mvrp applicant-state on page 248
- show mvrp dynamic-vlan-memberships on page 250
- show mvrp registration-state on page 252
- show mvrp statistics on page 254

List of Sample Output show mvrp interface on page 251

Output Fields Table 20 on page 251 lists the output fields for the **show mvrp interface** command. Output fields are listed in the approximate order in which they appear.

Table 20: show mvrp interface Output Fields

Field Name	Field Description
Interface	Interface on which MVRP is configured.
Status	Status of the MVRP: **Enabled** or **Disabled**.
Registration Mode	Registration for the interface: **Fixed**, **Forbidden**, or **Normal**.
Applicant Mode	Applicant mode.

Sample Output

show mvrp interface

```
user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface       Status      Registration      Applicant
                            Mode              Mode
ge-11/2/8       Enabled     Normal            Normal
ge-11/0/9       Enabled     Normal            Normal
ge-11/3/0       Enabled     Normal            Normal
```

show mvrp registration-state

Syntax show mvrp registration-state

Release Information Command introduced in Junos OS Release 10.1.

Description For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) registration state information.

Required Privilege Level view

Related Documentation

- show mvrp on page 246
- show mvrp dynamic-vlan-memberships on page 250
- show mvrp interface on page 251
- show mvrp statistics on page 254

List of Sample Output show mvrp registration-state on page 253

Output Fields Table 21 on page 252 lists the output fields for the **show mvrp registration-state** command. Output fields are listed in the approximate order in which they appear.

Table 21: show mvrp registration-state Output Fields

Field Name	Field Description
VLAN Id	Displays the Virtual LAN (VLAN) ID number.
Interface	Displays the interface number associated with the VLAN ID.
Registrar State	Displays whether the registrar state is **Registered** or **Empty**.
Forced State	Displays whether the forced state is **Registered** or **Empty**.
Managed State	Displays one of the following states: • **Fixed**—VLANs always stay in a registered state and are declared as such on all other forwarding ports. • **Normal**—VLANs participate in the MVRP protocol and honor incoming join requests normally. • **Forbidden**—VLANs ignore the incoming join requests and always stay in an unregistered state.
STP State	Displays whether the Spanning Tree Protocol (STP) is **Blocking** or **Forwarding**.

Sample Output

show mvrp registration-state

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'

VLAN Id   Interface   Registrar   Forced      Managed    STP
                      State       State       State      State
    100   ge-11/2/8   Empty       Registered  Fixed      Forwarding
          ge-11/0/9   Empty       Empty       Normal     Forwarding
          ge-11/3/0   Registered  Registered  Normal     Forwarding
    101   ge-11/2/8   Empty       Registered  Fixed      Forwarding
          ge-11/0/9   Empty       Empty       Normal     Forwarding
          ge-11/3/0   Registered  Registered  Normal     Forwarding
```

show mvrp statistics

Syntax show mvrp statistics

Release Information Command introduced in Junos OS Release 10.1.

Description For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.

Required Privilege Level view

Related Documentation

- show mvrp on page 246
- show mvrp applicant-state on page 248
- show mvrp dynamic-vlan-memberships on page 250
- show mvrp interface on page 251
- show mvrp registration-state on page 252

List of Sample Output show mvrp statistics on page 255

Output Fields Table 22 on page 254 lists the output fields for the **show mvrp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 22: show mvrp statistics Output Fields

Field Name	Field Description
interface name	Interface for which MVRP statistics are displayed.
VLAN IDs registered	Number of Virtual LAN (VLAN) IDs registered.
Sent MVRP PDUs	Number of MRPDU messages transmitted from the router.
Received MVRP PDUs without error	Number of MRPDU messages received on the router.
Received MVRP PDUs with error	Number of invalid MRPDU messages received on the router.

Sample Output

show mvrp statistics

```
user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name                 : ge-11/2/8
VLAN IDs registered            : 0
Sent MVRP PDUs                 : 1467
Received MVRP PDUs without error: 0
Received MVRP PDUs with error   : 0

Interface name                 : ge-11/0/9
VLAN IDs registered            : 0
Sent MVRP PDUs                 : 1418
Received MVRP PDUs without error: 702
Received MVRP PDUs with error   : 0

Interface name                 : ge-11/3/0
VLAN IDs registered            : 2
Sent MVRP PDUs                 : 1524
Received MVRP PDUs without error: 1366
Received MVRP PDUs with error   : 0
```

show vlans

Syntax

show vlans
<brief | detail | extensive>
<instance *instance-name*>
<logical-system *logical-system-name*>
<operational>
<*vlan-name*)>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

(MX Series routers and EX Series switches only) Display VLAN information.

Options

none—Display information for all VLANs.

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

instance *instance-name*—(Optional) Display information for the specified routing instance.

logical-system *logical-system-name*—(Optional) Display Ethernet-switching statistics information for the specified logical system.

operational—(Optional) Display information for the operational routing instances.

***vlan-name*)**— (Optional) Display information about the specified VLAN.

Required Privilege Level

view

List of Sample Output

[xref target has no title]
show vlans detail on page 257

Sample Output

```
Routing instance        VLAN name             Tag         Interfaces
VPLS-1                  __VPLS-1__               all
                                                            ae1.0
VPLS-2                  __VPLS-2__               all
                                                            ae3.0
                                                            ge-3/1/2.0
                                                            vt-3/3/10.1048576
default-switch          VLAN1000                 1000
                                                            ae26.0
default-switch          VLAN101                  101
                                                            ae20.0
default-switch          VLAN102                  102
                                                            ae20.0
default-switch          VLAN103                  103
                                                            ae20.0
default-switch          VLAN104                  104
                                                            ae20.0
default-switch          VLAN105                  105
                                                            ae20.0
default-switch          VLAN106                  106
                                                            ae20.0
default-switch          VLAN107                  107
                                                            ae20.0
default-switch          VLAN108                  108
                                                            ae20.0
[...output truncated...]
```

show vlans detail

```
user@host> show vlans detail
Routing instance: VPLS-1
  VLAN Name: __VPLS-1__                     State: Active
Tag: all
Internal index: 2, Generation Index:  , Origin: Dynamic
Interfaces:
    ae1.0,tagged
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0

Routing instance: VPLS-2
  VLAN Name: __VPLS-2__                     State: Active
Tag: all
Internal index: 3, Generation Index:  , Origin: Dynamic
Interfaces:
    ae3.0,tagged
    ge-3/1/2.0,tagged
    vt-3/3/10.1048576,tagged
Number of interfaces: Tagged 3    , Untagged 0
Total MAC count: 4

Routing instance: default-switch
  VLAN Name: VLAN1000                       State: Active
Tag: 1000
Internal index: 4, Generation Index: 1, Origin: Static
Layer 3 interface: irb.1000
Interfaces:
    ae26.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 0
```

```
Routing instance: default-switch
  VLAN Name: VLAN101                        State: Active
Tag: 101
Internal index: 5, Generation Index: 2, Origin: Static
Layer 3 interface: irb.101
Interfaces:
    ae20.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 1

Routing instance: default-switch
  VLAN Name: VLAN102                        State: Active
Tag: 102
Internal index: 6, Generation Index: 3, Origin: Static
Layer 3 interface: irb.102
Interfaces:
    ae20.0,tagged,trunk
Number of interfaces: Tagged 1    , Untagged 0
Total MAC count: 1
[...output truncated...]
```

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 260

Output Fields

Table 23 on page 259 lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

Table 23: 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 23: 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 maintenance point that is sending the linktrace message.
Next-hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.

Sample Output

traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:90:69:7e:01:ff
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:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:01:02:03:04:05
       4       60       00:01:02:03:04:05         00:00:00:00:00:00
```