



Junos[®] OS for EX Series Ethernet Switches

Interfaces for EX9200 Switches

Release
12.3



Published: 2013-04-03

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS for EX Series Ethernet Switches Interfaces for EX9200 Switches

Release 12.3

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	Aggregated Ethernet Overview	3
	Aggregated Ethernet Interfaces Overview	3
	Platform Support for Aggregated Ethernet Interfaces	3
	Configuration Guidelines for Aggregated Ethernet Interfaces	4
	Load Balancing and Ethernet Link Aggregation Overview	5
Chapter 2	IP-Directed Broadcast Overview	7
	Understanding Targeted Broadcast	7
Chapter 3	Reverse Path Forwarding	9
	Understanding Unicast Reverse Path Forwarding	9
	Understanding Multicast Reverse Path Forwarding	9
	RPF Table	10
Part 2	Configuration	
Chapter 4	Configuration	15
	Configuring a Layer 2 Virtual Switch	15
	Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port	16
	Understanding Layer 2 Virtual Switches Instances	17
	Configuring VLAN Encapsulation	17
	Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface	18
	Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface	18
	Rewriting the Inner and Outer VLAN Tags	19
	Rewriting the VLAN Tag on Tagged Frames	19

Binding VLAN IDs to Logical Interfaces	20
Binding VLAN IDs to Logical Interfaces Overview	21
Binding a VLAN ID to a Logical Interface	21
Binding a VLAN ID to a Single-Tag Logical Interface	21
Binding a VLAN ID to a Dual-Tag Logical Interface	22
Binding a Range of VLAN IDs to a Logical Interface	22
Binding a Range of VLAN IDs to a Single-Tag Logical Interface	22
Binding a Range of VLAN IDs to a Dual-Tag Logical Interface	22
Example: Binding Ranges VLAN IDs to Logical Interfaces	23
Binding a List of VLAN IDs to a Logical Interface	23
Binding a List of VLAN IDs to a Single-Tag Logical Interface	24
Binding a List of VLAN IDs to a Dual-Tag Logical Interface	24
Example: Binding Lists of VLAN IDs to Logical Interfaces	24
Configuring a Logical Interface for Access Mode	26
Configuring Junos OS for Supporting Aggregated Devices	26
Configuring Virtual Links for Aggregated Devices	26
Configuring LACP Link Protection at the Chassis Level	27
Enabling LACP Link Protection	28
Configuring System Priority	28
Configuring the Maximum Links Limit	28
Configuring an Aggregated Ethernet Interface	29
Deleting an Aggregated Ethernet Interface	30
Configuring the Number of Aggregated Ethernet Interfaces on the Device (Enhanced Layer 2 Software CLI Procedure)	30
Example: Configuring Aggregated Ethernet Interfaces	31
Configuring Tagged Aggregated Ethernet Interfaces	32
Configuring Untagged Aggregated Ethernet Interfaces	33
Configuring Aggregated Ethernet Minimum Links	33
Configuring Load Balancing on a LAG Link	34
Example: Configuring Load Balancing on a LAG Link	35
Configuring Multichassis Link Aggregation	35
Configuring Aggregated Ethernet LACP	37
Configuring the LACP Interval	38
Configuring LACP Link Protection	39
Configuring LACP System Priority	40
Configuring LACP System Identifier	40
Configuring LACP administrative Key	41
Configuring LACP Port Priority	41
Tracing LACP Operations	41
LACP Limitations	42
Example: Configuring Aggregated Ethernet LACP	42
Configuring Targeted Broadcast	43
Configuring Unicast RPF	44
Configuring Unicast RPF Strict Mode	45
Configuring Unicast RPF Loose Mode	46
Configuring Unicast RPF Loose Mode with Ability to Discard Packets	46

	Unicast RPF and Default Routes	47
	Unicast RPF Behavior with a Default Route	47
	Unicast RPF Behavior Without a Default Route	48
	Unicast RPF with Routing Asymmetry	48
	Configuring Unicast RPF on a VPN	49
	Example: Configuring Unicast RPF on a VPN	49
	Example: Configuring Unicast RPF	49
Chapter 5	Configuration Statements	51
	[edit chassis] Hierarchy Level	51
	[edit dynamic-profiles] Hierarchy Level	59
	[edit forwarding-options rpf-loose-mode-discard] Hierarchy Level	60
	[edit interfaces] Hierarchy Level	60
	[edit multi-chassis] Hierarchy Level	71
	[edit protocols isis] Hierarchy Level	71
	Layer 2 Routing Instances Configuration Hierarchy	74
	bandwidth (Interfaces)	77
	chassis-id	78
	filter	79
	flow-control	80
	forward-and-send-to-re	81
	forward-only	81
	gratuitous-arp-reply	82
	group (RPF Selection)	83
	l2-domain-id-for-l3	83
	lACP (Aggregated Ethernet)	84
	layer3-domain-identifier	85
	link-protection	86
	mc-ae-id	87
	mode (QFX Series)	87
	mode (Interfaces)	88
	multicast-rpf-routes	88
	next-hop (PIM RPF Selection)	89
	no-gratuitous-arp-request	89
	no-local-switching	90
	policer (MAC)	91
	prefix-list (PIM RPF Selection)	92
	redundancy-group	93
	rpf-check (Dynamic Profiles)	93
	rpf-check (interfaces)	94
	rpf-check-policy (Routing Options RPF)	95
	rpf-loose-mode-discard	96
	rpf-selection	97
	source (PIM RPF Selection)	98
	status-control	98
	targeted-broadcast	99
	traceoptions (Individual Interfaces)	100
	traps	101
	unicast-reverse-path	102

	unidirectional	103
	vlan-tagging	104
	wildcard-source (PIM RPF Selection)	105
Part 3	Administration	
Chapter 6	Routine Monitoring	109
	Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface	109
Chapter 7	Operational Commands	111
	Common Output Fields Description	111
	Destination Class Field	111
	Enabled Field	111
	Filters Field	112
	Flags Fields	112
	Addresses, Flags Field	112
	Device Flags Field	113
	Family Flags Field	113
	Interface Flags Field	114
	Link Flags Field	115
	Logical Interface Flags Field	115
	Label-Switched Interface Traffic Statistics Field	116
	Policer Field	117
	Protocol Field	117
	RPF Failures Field	118
	Source Class Field	118
	clear interfaces statistics	119
	show interfaces (10-Gigabit Ethernet)	120
	show interfaces (Discard)	147
	show interfaces (Gigabit Ethernet)	153
	show interfaces (Serial)	177
	show interfaces extensive	191
	show interfaces queue	211

List of Figures

Part 1	Overview	
Chapter 3	Reverse Path Forwarding	9
	Figure 1: Multicast Routers and the RPF Check	10
Part 2	Configuration	
Chapter 4	Configuration	15
	Figure 2: Unicast RPF with Routing Asymmetry	48

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 2	Configuration	
Chapter 4	Configuration	15
	Table 3: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces	21
Part 3	Administration	
Chapter 7	Operational Commands	111
	Table 4: show interfaces Gigabit Ethernet Output Fields	121
	Table 5: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . .	135
	Table 6: Discard show interfaces Output Fields	147
	Table 7: show interfaces Gigabit Ethernet Output Fields	154
	Table 8: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . .	167
	Table 9: show interfaces (Serial) Output Fields	177
	Table 10: Layer 2 Overhead, Transmitted Packets/Bytes	212
	Table 11: show interfaces queue Output Fields	215
	Table 12: Byte Count by PIC Type	218

About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

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.xml;
    }
  }
}
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.xml; }
```

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
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	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 <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> 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 <i>(string1 string2 string3)</i>
# (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]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; 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.	<ul style="list-style-type: none"> 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

- [Aggregated Ethernet Overview on page 3](#)
- [IP-Directed Broadcast Overview on page 7](#)
- [Reverse Path Forwarding on page 9](#)

CHAPTER 1

Aggregated Ethernet Overview

- [Aggregated Ethernet Interfaces Overview on page 3](#)
- [Load Balancing and Ethernet Link Aggregation Overview on page 5](#)

Aggregated Ethernet Interfaces Overview

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



NOTE: For information about configuring circuit cross-connects over aggregated Ethernet, see [Circuit and Translational Cross-Connects Overview](#).

Platform Support for Aggregated Ethernet Interfaces

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

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

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

All other routers support a maximum of 8 physical interfaces per aggregated Ethernet bundle.

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

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

Configuration Guidelines for Aggregated Ethernet Interfaces

Simple filters are not supported for interfaces in aggregated Ethernet bundles:

- On M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
- On MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

For more information about simple filters, see the Junos OS Class of Service Configuration Guide.

On the aggregated bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific capabilities, see Gigabit Ethernet Accounting and Policing Overview.

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

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

- **inner-tag-protocol-id**
- **inner-vlan-id**
- **pop-pop**
- **pop-swap**
- **push-push**
- **swap-push**
- **swap-swap**

In all cases, you must set the number of aggregated Ethernet interfaces on the chassis. You can also set the link speed and the minimum links in a bundle.

Related Documentation

- **inner-tag-protocol-id**
- **inner-vlan-id**

- **pop-pop**
- **pop-swap**
- **push-push**
- **swap-push**
- **swap-swap**
- Gigabit Ethernet Accounting and Policing Overview
- Junos® OS Ethernet Interfaces

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 480 LAG bundles on a Juniper Networks MX Series Ethernet Services Router. Each LAG bundle contains up to 16 links.

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. You can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement, see [“Configuring Load Balancing on a LAG Link” on page 34](#). In a Layer 2 switch, one link is overutilized and other links are underutilized.

Related Documentation

- [Configuring Load Balancing on a LAG Link on page 34](#)
- [Load Balancing on a LAG Link on page 35](#)
- **payload**

CHAPTER 2

IP-Directed Broadcast Overview

- [Understanding Targeted Broadcast on page 7](#)

Understanding Targeted Broadcast

Targeted broadcast is a process of flooding a target subnet with Layer 3 broadcast IP packets originating from a different subnet. The intent of targeted broadcast is to flood the target subnet with the broadcast packets on a LAN interface without broadcasting to the entire network. Targeted broadcast is configured with various options on the egress interface of the router or switch and the IP packets are broadcast only on the LAN (egress) interface. Targeted broadcast helps you implement remote administration tasks such as backups and wake-on LAN (WOL) on a LAN interface, and supports virtual routing and forwarding (VRF) instances.

Regular Layer 3 broadcast IP packets originating from a subnet are broadcast within the same subnet. When these IP packets reach a different subnet, they are forwarded to the Routing Engine (to be forwarded to other applications). Because of this, remote administration tasks such as backups cannot be performed on a particular subnet through another subnet. As a workaround you can enable targeted broadcast, to forward broadcast packets that originate from a different subnet.

Layer 3 broadcast IP packets have a destination IP address that is a valid broadcast address for the target subnet. These IP packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. In the destination subnet, if the receiving router has targeted broadcast enabled on the egress interface, the IP packets are forwarded to an egress interface and the Routing Engine or to an egress interface only. The IP packets are then translated into broadcast IP packets which flood the target subnet only through the LAN interface (if there is no LAN interface, the packets are discarded), and all hosts on the target subnet receive the IP packets. If targeted broadcast is not enabled on the receiving router, the IP packets are treated as regular Layer 3 broadcast IP packets and are forwarded to the Routing Engine. If targeted broadcast is enabled without any options, the IP packets are discarded.

Targeted broadcast can be configured to forward the IP packets only to an egress interface, which is helpful when the router is flooded with packets to process, or to both an egress interface and the Routing Engine.



NOTE: Any firewall filter that is configured on the Routing Engine loopback interface (lo0) cannot be applied to IP packets that are forwarded to the Routing Engine as a result of a targeted broadcast. This is because broadcast packets are forwarded as flood next hop and not as local next hop traffic, and you can only apply a firewall filter to local next hop routes for traffic directed towards the Routing Engine.

- Related Documentation**
- [Configuring Targeted Broadcast on page 43](#)
 - [targeted-broadcast on page 99](#)

CHAPTER 3

Reverse Path Forwarding

- [Understanding Unicast Reverse Path Forwarding on page 9](#)
- [Understanding Multicast Reverse Path Forwarding on page 9](#)

Understanding Unicast Reverse Path Forwarding

IP spoofing can occur during a denial-of-service (DoS) attack. IP spoofing allows an intruder to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination. This type of spoofing is harmful because it consumes the destination's resources.

A unicast reverse-path-forwarding (RPF) check is a tool to reduce forwarding of IP packets that might be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router or switch determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router or switch forwards the packet to the destination address. If it is not from a valid path, the router or switch discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family.



NOTE: Reverse path forwarding is not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.

Related Documentation

- [Example: Configuring Unicast Reverse-Path-Forwarding Check](#)

Understanding Multicast Reverse Path Forwarding

Unicast forwarding decisions are typically based on the destination address of the packet arriving at a router. The unicast routing table is organized by destination subnet and mainly set up to forward the packet toward the destination.

In multicast, the router or switch forwards the packet away from the source to make progress along the distribution tree and prevent routing loops. The router's or switch's multicast forwarding state runs more logically by organizing tables based on the reverse

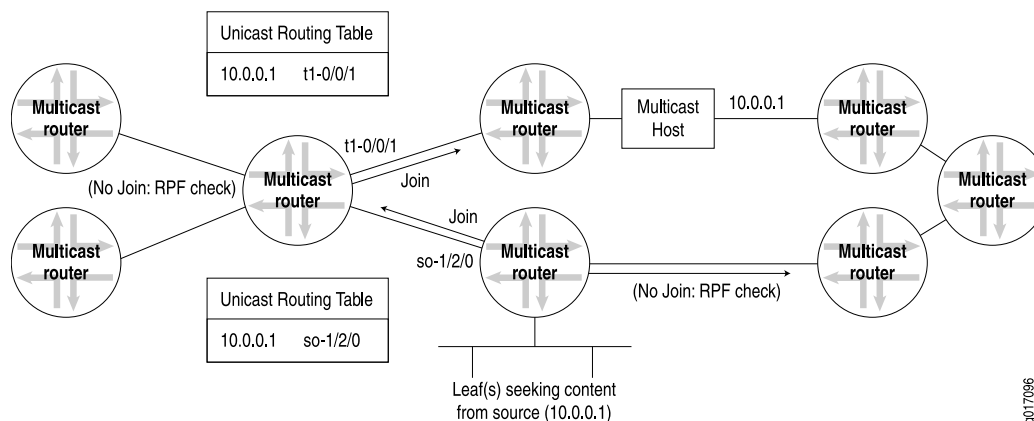
path, from the receiver back to the root of the distribution tree. This process is known as *reverse-path forwarding (RPF)*.

The router or switch adds a branch to a distribution tree depending on whether the request for traffic from a multicast group passes the reverse-path-forwarding check (RPF check). Every multicast packet received must pass an RPF check before it is eligible to be replicated or forwarded on any interface.

The RPF check is essential for every router's multicast implementation. When a multicast packet is received on an interface, the router or switch interprets the source address in the multicast IP packet as the destination address for a unicast IP packet. The source multicast address is found in the unicast routing table, and the outgoing interface is determined. If the outgoing interface found in the unicast routing table is the same as the interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped because the incoming interface is not on the *shortest path* back to the source.

Figure 1 on page 10 shows how multicast routers can use the unicast routing table to perform an RPF check and how the results obtained at each router determine where join messages are sent.

Figure 1: Multicast Routers and the RPF Check



Routers and switches can build and maintain separate tables for RPF purposes. The router must have some way to determine its RPF interface for the group, which is the interface topologically closest to the root. For greatest efficiency, the distribution tree follows the shortest-path tree topology. The RPF check helps to construct this tree.

RPF Table

The RPF table plays the key role in the multicast router or switch. The RPF table is consulted for every RPF check, which is performed at intervals on multicast packets entering the multicast router. Distribution trees of all types rely on the RPF table to form properly, and the multicast forwarding state also depends on the RPF table.

RPF checks are performed only on unicast addresses to find the upstream interface for the multicast source or RP.

The routing table used for RPF checks can be the same routing table used to forward unicast IP packets, or it can be a separate routing table used only for multicast RPF checks. In either case, the RPF table contains only unicast routes, because the RPF check is performed on the source address of the multicast packet, not the multicast group destination address, and a multicast address is forbidden from appearing in the source address field of an IP packet header. The unicast address can be used for RPF checks because there is only one source host for a particular stream of IP multicast content for a multicast group address, although the same content could be available from multiple sources.

If the same routing table used to forward unicast packets is also used for the RPF checks, the routing table is populated and maintained by the traditional unicast routing protocols such as BGP, IS-IS, OSPF, and the Routing Information Protocol (RIP). If a dedicated multicast RPF table is used, this table must be populated by some other method. Some multicast routing protocols (such as the Distance Vector Multicast Routing Protocol [DVMRP]) essentially duplicate the operation of a unicast routing protocol and populate a dedicated RPF table. Others, such as PIM, do not duplicate routing protocol functions and must rely on some other routing protocol to set up this table, which is why PIM is *protocol independent*.

Some traditional routing protocols such as BGP and IS-IS now have extensions to differentiate between different sets of routing information sent between routers and switches for unicast and multicast. For example, there is multiprotocol BGP (MBGP) and multitopology routing in IS-IS (M-IS-IS). IS-IS routes can be added to the RPF table even when special features such as traffic engineering and “shortcuts” are turned on. Multicast Open Shortest Path First (MOSPF) also extends OSPF for multicast use, but goes further than MBGP or M-IS-IS and makes MOSPF into a complete multicast routing protocol on its own. When these routing protocols are used, routes can be tagged as multicast RPF routers and used by the receiving router differently than the unicast routing information.

Using the main unicast routing table for RPF checks provides simplicity. A dedicated routing table for RPF checks allows a network administrator to set up separate paths and routing policies for unicast and multicast traffic, allowing the multicast network to function more independently of the unicast network.

By default, PIM uses **inet.0** as its reverse-path forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and for the RP address. PIM can optionally use **inet.2** as its RPF routing table group. The **inet.2** routing table is organized more efficiently for RPF checks.

Once configured, the RPF routing table must be applied to a PIM as a routing table group.

PART 2

Configuration

- [Configuration on page 15](#)
- [Configuration Statements on page 51](#)

CHAPTER 4

Configuration

- [Configuring a Layer 2 Virtual Switch on page 15](#)
- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 16](#)
- [Understanding Layer 2 Virtual Switches Instances on page 17](#)
- [Configuring VLAN Encapsulation on page 17](#)
- [Rewriting the Inner and Outer VLAN Tags on page 19](#)
- [Rewriting the VLAN Tag on Tagged Frames on page 19](#)
- [Binding VLAN IDs to Logical Interfaces on page 20](#)
- [Configuring a Logical Interface for Access Mode on page 26](#)
- [Configuring Junos OS for Supporting Aggregated Devices on page 26](#)
- [Configuring an Aggregated Ethernet Interface on page 29](#)
- [Deleting an Aggregated Ethernet Interface on page 30](#)
- [Configuring the Number of Aggregated Ethernet Interfaces on the Device \(Enhanced Layer 2 Software CLI Procedure\) on page 30](#)
- [Example: Configuring Aggregated Ethernet Interfaces on page 31](#)
- [Configuring Tagged Aggregated Ethernet Interfaces on page 32](#)
- [Configuring Untagged Aggregated Ethernet Interfaces on page 33](#)
- [Configuring Aggregated Ethernet Minimum Links on page 33](#)
- [Configuring Load Balancing on a LAG Link on page 34](#)
- [Example: Configuring Load Balancing on a LAG Link on page 35](#)
- [Configuring Multichassis Link Aggregation on page 35](#)
- [Configuring Aggregated Ethernet LACP on page 37](#)
- [Configuring Targeted Broadcast on page 43](#)
- [Configuring Unicast RPF on page 44](#)

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 16](#)

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 15](#)

Understanding Layer 2 Virtual Switches Instances

At Layer 2, you can group one or more VLANs into a single routing instance to form a virtual switch instance. A virtual switch instance is composed of VLANs. The virtual switch instance isolates a LAN segment and contains most Layer 2 functions, such as spanning-tree protocol instances and VLAN ID spaces, into its own smaller, logical network. Splitting Layer 2 traffic using virtual switch instances allows you to more logically organize your Layer 2 traffic into multiple “virtual” Layer 2 networks.

A default virtual switch, called default-switch, is automatically created when a virtual switch is configured. All Layer 2 traffic not assigned to a VLAN in a virtual switch automatically becomes part of the default virtual switch.

You can configure a virtual switch to participate only in Layer 2 bridging and optionally to perform Layer 3 routing. In addition, you can configure spanning-tree protocols (STPs) within the virtual switch to prevent forwarding loops. For more information about how to configure Layer 2 logical ports on an interface, see the Junos® OS Network Interfaces.

You can associate one or more logical interfaces configured as trunk interfaces with a virtual switch. A trunk interface, or Layer 2 trunk port, enables you to configure a logical interface to represent multiple VLANs on the physical interface. For more information about how to configure trunk interfaces, see the Junos® OS Network Interfaces.

You can also configure Layer 2 forwarding and learning properties for the virtual switch.

Related Documentation

- [Configuring a Layer 2 Virtual Switch on page 15](#)
- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 16](#)
- [Configuring a Layer 2 Control Protocol Routing Instance](#)

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

Rewriting the Inner and Outer VLAN Tags

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

```
swap-swap;
```

You can include this statement at the following hierarchy levels:

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

See 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
 - output-vlan-map
 - swap-swap
 - unit
 - Junos® OS Ethernet Interfaces

Rewriting the VLAN Tag on Tagged Frames

To rewrite the VLAN tag on all tagged frames entering the interface to a specified VLAN ID and TPID, include the **swap**, **tag-protocol-id**, and **vlan-id** statements in the input VLAN map:

```
input-vlan-map {
  swap;
  vlan-id number;
  tag-protocol-id tpid;
```

```
}
```

To rewrite the VLAN tag on all tagged frames exiting the interface to a specified VLAN ID and TPID, include the **swap** and **tag-protocol-id** statements in the output VLAN map:

```
output-vlan-map {  
    swap;  
    vlan-id number;  
    tag-protocol-id tpid;  
}
```

You can include these statements at the following hierarchy levels:

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

You cannot include both the **swap** statement and the **vlan-id** statement in the output VLAN map configuration. If you include the **swap** statement in the configuration, the VLAN ID in outgoing frames is rewritten to the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see 802.1Q VLANs Overview.

The swap operation works on the outer tag only, whether or not you include the **stacked-vlan-tagging** statement in the configuration. For more information, see Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags.

Related Documentation

- input-vlan-map
- output-vlan-map
- swap
- vlan-id
- tag-protocol-id
- unit
- For more information about binding a VLAN ID to the logical interface, see 802.1Q VLANs Overview.
- For more information about the swap operation, see Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags.
- Junos® OS Ethernet Interfaces

Binding VLAN IDs to Logical Interfaces

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

- [Binding VLAN IDs to Logical Interfaces Overview on page 21](#)
- [Binding a VLAN ID to a Logical Interface on page 21](#)

- [Binding a Range of VLAN IDs to a Logical Interface on page 22](#)
- [Binding a List of VLAN IDs to a Logical Interface on page 23](#)

Binding VLAN IDs to Logical Interfaces Overview

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, a range of VLAN IDs, or a list of VLAN IDs to the logical interface. [Table 3 on page 21](#) lists the configuration statements you use to bind VLAN IDs to logical interfaces, organized by scope of the VLAN IDs used to match incoming packets:

Table 3: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces

Scope of VLAN ID Matching	Type of VLAN Framing Supported on the Logical Interface	
	Single-Tag Framing	Dual-Tag Framing
VLAN ID	<code>vlan-id <i>vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid</i>.<<i>vlan-id</i>> inner <i>tpid</i><i>vlan-id</i>;</code>
VLAN ID Range	<code>vlan-id-range <i>vlan-id</i>–<i>vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid</i>.<i>vlan-id</i> inner-range <i>tpid</i>.<i>vlan-id</i>–<i>vlan-id</i>;</code>
VLAN ID List	<code>vlan-id-list [<i>vlan-id</i> <i>vlan-id</i>–<i>vlan-id</i>];</code>	<code>vlan-tags outer <<i>tpid</i>.><i>vlan-id</i> inner-list [<i>vlan-id</i> <i>vlan-id</i>–<i>vlan-id</i>];</code>

You can include all of the statements at the following hierarchy levels:

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



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

Binding a VLAN ID to a Logical Interface

A logical interface that you have associated (bound) to a particular VLAN ID will receive and forward incoming frames that contain a matching VLAN ID.

Binding a VLAN ID to a Single-Tag Logical Interface

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

```
vlan-id vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a VLAN ID to a Dual-Tag Logical Interface

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

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

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a Range of VLAN IDs to a Logical Interface

A VLAN range can be used by service providers to interconnect multiple VLANs belonging to a particular customer over multiple sites. Using a VLAN ID range conserves switch resources and simplifies configuration.

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

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

```
vlan-id-range vlan-id-vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

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

To bind a range of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner-list** option to specify the VLAN IDs as an inclusive range by separating the starting VLAN ID and ending VLAN ID with a hyphen.

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

You can include the statement at the following hierarchy levels:

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

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Example: Binding Ranges VLAN IDs to Logical Interfaces

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

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

Binding a List of VLAN IDs to a Logical Interface

In Junos OS Release 9.5 and later, on MX Series routers and in Junos OS Release 12.2R2 and later on EX Series switches, you can bind a list of VLAN IDs to a single logical interface, eliminating the need to configure a separate logical interface for every VLAN or VLAN range. A logical interface that accepts packets tagged with any VLAN ID specified in a VLAN ID list is called a *VLAN-bundled* logical interface.

You can use VLAN-bundled logical interfaces to configure circuit cross-connects between Layer 2 VPN routing instances or Layer 2 circuits. Using VLAN-bundled logical interfaces simplifies configuration and reduces use of system resources such as logical interfaces, next hops, and circuits.

As an alternative to configuring multiple logical interfaces (one for each VLAN ID and one for each range of VLAN IDs), you can configure a single VLAN-bundled logical interface based on a list of VLAN IDs.



NOTE: The `vlan-id` option is not supported to achieve VLAN normalization on VPLS instances that are configured with `vlan-id-list`. However, you can use the `vlan-maps` option to achieve VLAN normalization.

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

To bind a list of VLAN IDs to a single-tag logical interface, include the **vlan-id-list** statement. Specify the VLAN IDs in the list individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

```
vlan-id-list [ vlan-id vlan-id-vlan-id ];
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

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

To bind a list of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner-list** option to specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both:

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



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

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Example: Binding Lists of VLAN IDs to Logical Interfaces

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

```
[edit interfaces]
ge-1/1/0 {
  vlan-tagging; # Only for single-tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
```



```

        vlan-id-list [20 30–40 45];
    }
}
ge-1/1/1 {
    flexible-vlan-tagging; # Only for mixed tagging
    encapsulation flexible-ethernet-services;
    unit 10 {
        encapsulation vlan-ccc;
        vlan-id-list [1 10 20 30–40];
    }
    unit 20 {
        encapsulation vlan-ccc;
        vlan-tags outer 200 inner-list [50–60 80 90–100];
    }
}

```

In the example configuration above, **ge-1/1/0** supports single-tag logical interfaces, and **ge-1/1/1** supports mixed tagging. The single-tag logical interfaces **ge-1/1/0.10** and **ge-1/1/1.20** each bundle lists of VLAN IDs. The dual-tag logical interface **ge-1/1/1.20** bundles lists of inner VLAN IDs.



TIP: You can group a range of identical interfaces into an interface range and then apply a common configuration to that interface range. For example, in the above example configuration, both interfaces **ge-1/1/0** and **ge-1/1/1** have the same physical encapsulation type of **flexible-ethernet-services**. Thus you can define an interface range with the interfaces **ge-1/1/0** and **ge-1/1/1** as its members and apply the encapsulation type **flexible-ethernet-services** to that defined interface range. For more information about interface ranges, see [Configuring Interface Ranges](#).

- Related Documentation**
- [802.1Q VLANs Overview](#)
 - [Configuring Interface Ranges](#)
 - [Junos® OS Ethernet Interfaces](#)

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](#)

Configuring Junos OS for Supporting Aggregated Devices

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

Tasks for configuring aggregated devices are:

- [Configuring Virtual Links for Aggregated Devices on page 26](#)
- [Configuring LACP Link Protection at the Chassis Level on page 27](#)
- [Enabling LACP Link Protection on page 28](#)
- [Configuring System Priority on page 28](#)
- [Configuring the Maximum Links Limit on page 28](#)

Configuring Virtual Links for Aggregated Devices

To define virtual links, you need to specify the associations between physical and logical devices within the `[edit interfaces]` hierarchy, and assign the correct number of logical

devices by including the **device-count** statement at the **[edit chassis aggregated-devices ethernet]** and **[edit chassis aggregated-devices sonet]** hierarchy levels:

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

The maximum number of Ethernet logical interfaces that you can configure is 128. On M Series and T Series routers, you can configure a maximum number of 128 aggregated interfaces. On MX Series routers, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces are numbered from **ae0** through **ae127** for M Series and T Series routers, and the aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers. The maximum number of SONET/SDH logical interfaces is 16. The aggregated SONET/SDH interfaces are numbered from **as0** through **as15**.

Configuring LACP Link Protection at the Chassis Level

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

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

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



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

Enabling LACP Link Protection

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

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

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



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

Configuring System Priority

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

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

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

Configuring the Maximum Links Limit

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

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

maximum-links *maximum-links-limit*;

Related Documentation

- [Configuring an Aggregated Ethernet Interface](#)
- [Junos® OS Ethernet Interfaces](#)
- [Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Switches](#)

Configuring an Aggregated Ethernet Interface

You can associate a physical interface with an aggregated Ethernet interface.

To configure an aggregated Ethernet interface:

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

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

2. Configure the aggregated Ethernet interface.

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

You specify the interface instance number *x* to complete the link association; *x* can be from 0 through 480, for a total of 480 aggregated interfaces on MX Series routers or EX9200 switches. You must also include a statement defining **aex** at the **[edit interfaces]** hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see Ethernet Interfaces Overview.



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

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

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



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

Related Documentation

- [Configuring an Aggregated Ethernet Interface on page 29](#)
- [Configuring the Number of Aggregated Ethernet Interfaces on the Device \(Enhanced Layer 2 Software CLI Procedure\) on page 30](#)
- [Deleting an Aggregated Ethernet Interface on page 30](#)

- [Aggregated Ethernet Interfaces Overview on page 3](#)
- Junos® OS Ethernet Interfaces

Deleting an Aggregated Ethernet Interface

There are two approaches to deleting an aggregated Ethernet interface:

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

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

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

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

2. Delete the interface from the device count.

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

Related Documentation

- [Configuring an Aggregated Ethernet Interface](#)
- [Configuring the Number of Aggregated Ethernet Interfaces on the Device](#)
- [Aggregated Ethernet Interfaces Overview on page 3](#)
- Junos® OS Ethernet Interfaces

Configuring the Number of Aggregated Ethernet Interfaces on the Device (Enhanced Layer 2 Software CLI Procedure)

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them.

On MX Series routers and EX9200 switches, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers and EX9200 switches.

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

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

2. Set the number of aggregated Ethernet interfaces.

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

```
user@host# set device-count number
```

You must also specify the constituent physical links by including the **802.3ad** statement at the **[edit interfaces *interface-name* ether-options]** or **[edit interfaces *interface-name* ether-options]** hierarchy level.

Related Documentation

- For information about physical links, see [Configuring an Aggregated Ethernet Interface on page 29](#)
- Junos® OS Ethernet Interfaces
- For information about configuring aggregated devices, see the Junos OS System Basics Configuration Guide.

Example: Configuring Aggregated Ethernet Interfaces

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

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

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

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

- Related Documentation**
- [Junos® OS Ethernet Interfaces](#)
 - [Configure 'link-speed' for Gigabit Ethernet based Aggregate Ethernet interface bundles](#)

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces aex]** hierarchy level:

```
[edit interfaces aex]  
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

You can include this statement at the following hierarchy levels:

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

For more information about the **vlan-tagging** and **vlan-id** statements, see 802.1Q VLANs Overview.

- Related Documentation**
- [vlan-id](#)
 - [vlan-tagging on page 104](#)

Configuring Untagged Aggregated Ethernet Interfaces

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

- You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.
- You cannot include the **vlan-id** statement in the configuration of the logical interface.

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

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

Related Documentation

- For more information about configuring LACP, see [Configuring Aggregated Ethernet LACP on page 37](#).
- Junos® OS Ethernet Interfaces

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```
2. Configure the minimum number of links.

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

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links *number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

- Related Documentation**
- [aggregated-ether-options](#)
 - [minimum-links](#)
 - [Junos® OS Ethernet Interfaces](#)

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.



NOTE: For Dense Port Concentrators (DPC), any change in the hash key configuration requires a reboot for the changes to take effect. For Modular Port Concentrators (MPC) the reboot is not required.

For more information about link aggregation group (LAG) configuration, see the Junos® OS Network Interfaces.

- Related Documentation**
- [Load Balancing and Ethernet Link Aggregation on page 5](#)
 - [Load Balancing on a LAG Link on page 35](#)

Example: Configuring Load Balancing on a LAG Link

This example configures the load-balancing hash key to use the source Layer 3 IP address option and Layer 4 header fields as well as the source and destination MAC addresses for load balancing on a link aggregation group (LAG) link:

```
[edit]
forwarding-options {
  hash-key {
    family multiservice {
      source-mac;
      destination-mac;
      payload {
        ip {
          layer-3 {
            source-ip-only;
          }
          layer-4;
        }
      }
    }
  }
}
```

Related Documentation

- [Load Balancing and Ethernet Link Aggregation on page 5](#)
- [Configuring Load Balancing on a LAG Link on page 34](#)

Configuring Multichassis Link Aggregation

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

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

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

MC-LAG includes the following functionality:

- Active standby mode is supported using Link Aggregation Control Protocol (LACP)
- MC-LAG operates only between two chassis.

- Layer 2 circuit functions are supported with **ether-ccc** encapsulation.
- VPLS functions are supported with **ether-vpls** and **vlan-vpls**.



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

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

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



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

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

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

- Related Documentation**
- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview
 - Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation
 - show interfaces mc-ae
 - Junos® OS Ethernet Interfaces

Configuring Aggregated Ethernet LACP

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

For Multichassis Link Aggregation (MC-LAG), you must specify the **system-id** and **admin key**. MC-LAG peers use the same **system-id** while sending the LACP messages. The **system-id** can be configured on the MC-LAG network device and synchronized between peers for validation.

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

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

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

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

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

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

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



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

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

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

Starting with Junos OS release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.

To override the IEEE 802.3ad standard and facilitate subsecond failover, include the **fast-failover** statement at the **[edit interfaces *interface-name* aggregated-ether-options lACP]** hierarchy level.

For more information, see the following sections:

- [Configuring the LACP Interval on page 38](#)
- [Configuring LACP Link Protection on page 39](#)
- [Configuring LACP System Priority on page 40](#)
- [Configuring LACP System Identifier on page 40](#)
- [Configuring LACP administrative Key on page 41](#)
- [Configuring LACP Port Priority on page 41](#)
- [Tracing LACP Operations on page 41](#)
- [LACP Limitations on page 42](#)
- [Example: Configuring Aggregated Ethernet LACP on page 42](#)

Configuring the LACP Interval

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lACP]** hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options lACP]
periodic interval;
```

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.



NOTE: Source address filtering does not work when LACP is enabled.

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the Routing Policy Configuration Guide.

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

Configuring LACP Link Protection



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

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the **link-protection** and **system-priority** statements. Configuring values at this level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

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



NOTE: LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

To enable LACP link protection for an aggregated Ethernet interfaces, use the **link-protection** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable;
  revertive;
  non-revertive;
}
```

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

If LACP link protection is configured to be nonrevertive at the global (**[edit chassis]** hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.



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

We strongly recommend you to use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces to any other vendor device. Otherwise, the vendor device (say a Layer 2 switch, or a router), will not be able to manage the traffic coming from the two link aggregated Ethernet bundle. As a result, you might observe the vendor device sending back the traffic to the backup member link of the aggregated Ethernet interface.

Currently, MX-MPC2-3D, MX-MPC2-3D-Q, MX-MPC2-3D-EQ, MX-MPC1-3D, MX-MPC1-3D-Q, and MPC-3D-16XGE-SFP do not drop traffic coming back to the backup link, whereas DPCE-R-Q-20GE-2XGE, DPCE-R-Q-20GE-SFP, DPCE-R-Q-40GE-SFP, DPCE-R-Q-4XGE-XFP, DPCE-X-Q-40GE-SFP, and DPCE-X-Q-4XGE-XFP drop traffic coming to the backup link.

Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-priority;
```

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

Configuring LACP System Identifier

To configure the LACP system identifier for aggregated Ethernet interfaces, use the **system-id** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-id system-id;
```

The user-defined system identifier in LACP enables two ports from two separate devices to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Configuring LACP administrative Key

To configure an administrative key for LACP, include the **admin-key *number*** statement at the **edit interfaces *ae x* aggregated-ether-options *lACP*** hierarchy level:

```
[edit interfaces ae x aggregated-ether-options-lACP]
admin-key number;
```



NOTE: You must configure MC-LAG to configure the **admin-key** statement. For more information about MC-LAG, see [“Configuring Multichassis Link Aggregation”](#) on page 35.

Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the **port-priority** statement at the **[edit interfaces *interface-name* ether-options 802.3ad aeX *lACP*]** or **[edit interfaces *interface-name* ether-options 802.3ad aeX *lACP*]** hierarchy levels:

```
[edit interfaces interface-name ether-options 802.3ad aeX lACP]
port-priority priority;
```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and are assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.



NOTE: Port aggregation selection (discussed above) is performed for the active link when LACP link protection is enabled. Without LACP link protection, port priority is not used in port aggregation selection.

Tracing LACP Operations

To trace the operations of the LACP process, include the **traceoptions** statement at the **[edit protocols *lACP*]** hierarchy level:

```
[edit protocols lACP]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

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

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

For general information about tracing, see the tracing and logging information in the Junos OS System Basics Configuration Guide.

LACP Limitations

LACP can link together multiple different physical interfaces, but only features that are supported across all of the linked devices will be supported in the resulting link aggregation group (LAG) bundle. For example, different PICs can support a different number of forwarding classes. If you use link aggregation to link together the ports of a PIC that supports up to 16 forwarding classes with a PIC that supports up to 8 forwarding classes, the resulting LAG bundle will only support up to 8 forwarding classes. Similarly, linking together a PIC that supports WRED with a PIC that does not support it will result in a LAG bundle that does not support WRED.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

LACP with VLAN-Tagged Aggregated Ethernet	<pre>[edit interfaces] ge--1/1/1 { ether-options { 802.3ad ae0; } } ae0 { aggregated-ether-options { lacp { active; } } vlan-tagging; unit 0 { vlan-id 100; family inet { address 10.1.1.2/24 { vrrp-group 0 { virtual-address 10.1.1.4; priority 200; } } } } }</pre>
--	--

```

    }
  }
}

```

Configure aggregated Ethernet LACP over an untagged interface:

LACP with Untagged Aggregated Ethernet

```

[edit interfaces]
ge-1/1/1 {
  ether-options-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}

```

- Related Documentation**
- [lacp on page 84](#)
 - [link-protection on page 86](#)
 - traceoptions
 - Junos® OS Ethernet Interfaces

Configuring Targeted Broadcast

You can configure targeted broadcast with different options to forward the IP packets destined for a Layer 3 broadcast address to an egress interface and the Routing Engine or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.

To enable targeted broadcast:

1. Configure the physical interface:

```

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

```

2. Configure the logical unit number:

```

[edit interfaces interface-name]

```

```
user@host# edit unit logical-unit-number
```

3. Configure the protocol family `inet`:

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

```
user@host# edit family inet
```

4. Configure targeted broadcast:

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

```
user@host# edit targeted-broadcast
```

5. Specify one of the following options:

- To send packets to the egress interface and to the Routing Engine:

```
[edit interfaces interface-name unit logical-unit-number family inet  
targeted-broadcast]
```

```
user@host# set forward-and-send-to-re
```

- To send packets to only the egress interface:

```
[edit interfaces interface-name unit logical-unit-number family inet  
targeted-broadcast]
```

```
user@host# set forward-only
```

6. Verify the configuration. The following example configures targeted broadcast to both the egress interface and the Routing Engine:

```
[edit interfaces interface-name unit logical-unit-number family inet targeted-broadcast]
```

```
user@host# up
```

```
user@host# show
```

```
targeted-broadcast {  
  forward-and-send-to-re;  
}
```

- Related Documentation
- [targeted-broadcast on page 99](#)
 - [Understanding Targeted Broadcast on page 7](#)

Configuring Unicast RPF

For interfaces that carry IPv4 or IPv6 traffic, you can reduce the impact of denial of service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.



NOTE: If you want to configure unicast RPF, your router must be equipped with the Internet Processor II application-specific integrated circuit (ASIC).

If you enable unicast RPF on live traffic, some packets are dropped while the packet forwarding components are updating.

For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.

The following sections describe unicast RPF in detail:

- [Configuring Unicast RPF Strict Mode on page 45](#)
- [Configuring Unicast RPF Loose Mode on page 46](#)
- [Unicast RPF and Default Routes on page 47](#)
- [Unicast RPF with Routing Asymmetry on page 48](#)
- [Configuring Unicast RPF on a VPN on page 49](#)
- [Example: Configuring Unicast RPF on page 49](#)

Configuring Unicast RPF Strict Mode

In strict mode, unicast RPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

If the incoming packet fails the unicast RPF check, the packet is not accepted on the interface. When a packet is not accepted on an interface, unicast RPF counts the packet and sends it to an optional fail filter. If the fail filter is not configured, the default action is to silently discard the packet.

The optional fail filter allows you to apply a filter to packets that fail the unicast RPF check. You can define the fail filter to perform any filter operation, including accepting, rejecting, logging, sampling, or policing.

When unicast RPF is enabled on an interface, Bootstrap Protocol (BOOTP) packets and Dynamic Host Configuration Protocol (DHCP) packets are not accepted on the interface. To allow the interface to accept BOOTP packets and DHCP packets, you must apply a fail filter that accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**. For a configuration example, see [“Example: Configuring Unicast RPF” on page 49](#).

For more information about unicast RPF, see the Junos OS Routing Protocols Configuration Guide. For more information about defining fail filters, see the Routing Policy Configuration Guide.

To configure unicast RPF, include the **rpf-check** statement:

```
rpf-check <fail-filter filter-name>;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]

Using unicast RPF can have several consequences when implemented with traffic filters:

- RPF fail filters are evaluated after input filters and before output filters.
- If you configure a filter counter for packets dropped by an input filter, and you want to know the total number of packets dropped, you must also configure a filter counter for packets dropped by the RPF check.
- To count packets that fail the RPF check and are accepted by the RPF fail filter, you must configure a filter counter.
- If an input filter forwards packets anywhere other than the **inet.0** or **inet6.0** routing tables, the unicast RPF check is not performed.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

Configuring Unicast RPF Loose Mode

By default, unicast RPF uses strict mode. Unicast RPF loose mode is similar to unicast RPF strict mode and has the same configuration restrictions. The only check in loose mode is whether the packet has a source address with a corresponding prefix in the routing table; loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. If a corresponding prefix is not found, unicast RPF loose mode does not accept the packet. As in strict mode, loose mode counts the failed packet and optionally forwards it to a fail filter, which either accepts, rejects, logs, samples, or polices the packet.

To configure unicast RPF loose mode, include the **mode**:

```
mode loose;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter filter-name>]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter filter-name>]

Configuring Unicast RPF Loose Mode with Ability to Discard Packets

Starting with Junos OS Release 12.1, unicast RPF loose mode has the ability to discard packets with the source address pointing to the discard interface. This feature is supported

on MX Series routers and on T Series routers with Type 1 FPCs, Type 2 FPCs, and Type 3 FPCs. Using unicast RPF loose mode, along with Remote Triggered Black Hole (RTBH) filtering, provides an efficient way to discard packets coming from known attack sources. BGP policies in edge routers ensure that packets with untrusted source addresses have their next hop set to a discard route. When a packet arrives at the router with an untrusted source address, unicast RPF performs a route lookup of the source address. Because the source address route points to a discard next hop, the packet is dropped and a counter is incremented. This feature is supported on both IPv4 (**inet**) and IPv6 (**inet6**) address families.

To configure unicast RPF loose mode with the ability to discard packets, include the **rpf-loose-mode-discard family inet** statement at the **[edit forwarding-options]** hierarchy level:

```
rpf-loose-mode-discard {
  family {
    inet;
  }
}
```

Unicast RPF and Default Routes

When the active route cannot be chosen from the routes in a routing table, the router chooses a default route. A default route is equivalent to an IP address of 0.0.0.0/0. If you configure a default route, and you configure unicast RPF on an interface that the default route uses, unicast RPF behaves differently than it does otherwise. For information about configuring default routes, see the Junos OS Routing Protocols Configuration Guide.

To determine whether the default route uses an interface, enter the **show route** command:

```
user@host> show route address
```

address is the next-hop address of the configured default route. The default route uses the interfaces shown in the output of the **show route** command.

The following sections describe how unicast RPF behaves when a default route uses an interface and when a default route does not use an interface:

- [Unicast RPF Behavior with a Default Route on page 47](#)
- [Unicast RPF Behavior Without a Default Route on page 48](#)

Unicast RPF Behavior with a Default Route

If you configure a default route that uses an interface configured with unicast RPF, unicast RPF behaves as follows:

- Loose mode—All packets are automatically accepted. For this reason, we recommend that you not configure unicast RPF loose mode on interfaces that the default route uses.
- Strict mode—The packet is accepted when either of the following is true:
 - The source address of the packet matches any of the routes (either default or learned) that can be originated from the interface. Note that routes can have multiple

destinations associated with them; therefore, if one of the destinations matches the incoming interface of the packet, the packet is accepted.

- The source address of the packet does not match any of the routes.

The packet is not accepted when either of the following is true:

- The source address of the packet does not match a prefix in the routing table.
- The interface does not expect to receive a packet with this source address prefix.

Unicast RPF Behavior Without a Default Route


If you do not configure a default route, or if the default route does not use an interface configured with unicast RPF, unicast RPF behaves as described in “[Configuring Unicast RPF Strict Mode](#)” on page 45 and “[Configuring Unicast RPF Loose Mode](#)” on page 46. To summarize, unicast RPF without a default route behaves as follows:

- Strict mode—The packet is not accepted when either of the following is true:
 - The packet has a source address that does not match a prefix in the routing table.
 - The interface does not expect to receive a packet with this source address prefix.
- Loose mode—The packet is not accepted when the packet has a source address that does not match a prefix in the routing table.

Unicast RPF with Routing Asymmetry

In general, we recommend that you not enable unicast RPF on interfaces that are internal to the network because internal interfaces are likely to have *routing asymmetry*. Routing asymmetry means that a packet’s outgoing and return paths are different. Routers in the core of the network are more likely to have asymmetric reverse paths than routers at the customer or provider edge. [Figure 2 on page 48](#) shows unicast RPF in an environment with routing asymmetry.

Figure 2: Unicast RPF with Routing Asymmetry



In [Figure 2 on page 48](#), if you enable unicast RPF on interface **so-0/0/0**, traffic destined for Router A is not rejected. If you enable unicast RPF on interface **so-1/0/1**, traffic from Router A is rejected.

If you need to enable unicast RPF in an asymmetric routing environment, you can use fail filters to allow the router to accept incoming packets that are known to be arriving by specific paths. For an example of a fail filter that accepts packets with a specific source and destination address, see “[Example: Configuring Unicast RPF](#)” on page 49.

Configuring Unicast RPF on a VPN

You can configure unicast RPF on a VPN interface by enabling unicast RPF on the interface and including the **interface** statement at the **[edit routing-instances routing-instance-name]** hierarchy level.

You can configure unicast RPF only on the interfaces you specify in the routing instance. This means the following:

- For Layer 3 VPNs, unicast RPF is supported on the CE router interface.
- Unicast RPF is not supported on core-facing interfaces.
- For virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

For more information about VPNs and virtual-router routing instances, see the Junos OS VPNs Configuration Guide. For more information about FBF, see the Junos OS Routing Protocols Configuration Guide.

Example: Configuring Unicast RPF on a VPN

Configure unicast RPF on a Layer 3 VPN interface:

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
}
[edit routing-instance]
VPN-A {
  interface so-0/0/0.0;
}
```

Example: Configuring Unicast RPF

Configure unicast RPF strict mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
    }
    address {
      255.255.255.255/32;
    }
  }
}
```

```
    }  
  }  
  then {  
    count rpf-dhcp-bootp-traffic;  
    accept;  
  }  
}  
term default {  
  then {  
    log;  
    reject;  
  }  
}  
}  
[edit]  
interfaces {  
  so-0/0/0 {  
    unit 0 {  
      family inet {  
        rpf-check fail-filter rpf-special-case-dhcp-bootp;  
      }  
    }  
  }  
}
```

- Related Documentation**
- [unicast-reverse-path on page 102](#)
 - Example: Configuring Unicast Reverse-Path-Forwarding Check

CHAPTER 5

Configuration Statements

- [\[edit chassis\] Hierarchy Level on page 51](#)
- [\[edit dynamic-profiles\] Hierarchy Level on page 59](#)
- [\[edit forwarding-options rpf-loose-mode-discard\] Hierarchy Level on page 60](#)
- [\[edit interfaces\] Hierarchy Level on page 60](#)
- [\[edit multi-chassis\] Hierarchy Level on page 71](#)
- [\[edit protocols isis\] Hierarchy Level on page 71](#)
- [Layer 2 Routing Instances Configuration Hierarchy on page 74](#)

[\[edit chassis\] Hierarchy Level](#)

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection {
          non-revertive;
        }
        system-priority;
      }
    }
    sonet {
      device-count number;
    }
    maximum-links maximum-links-limit;
  }
  alarm {
    ds1 {
      ais (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
    ethernet {
      link-down (ignore | red | yellow);
    }
    integrated-services {
      failure (ignore | red | yellow);
    }
    management-ethernet {
```

```

    link-down (ignore | red | yellow);
}
relay
input {
    port port-number {
        mode (close | open);
        trigger (ignore | red | yellow);
    }
}
output {
    port port-number {
        input-relay input-relay;
        mode (close | open);
        temperature;
    }
}
serial {
    cts-absent (ignore | red | yellow);
    dcd-absent (ignore | red | yellow);
    dsr-absent (ignore | red | yellow);
    loss-of-rx-clock (ignore | red | yellow);
    loss-of-tx-clock (ignore | red | yellow);
    tm-absent (ignore | red | yellow);
}
services {
    hw-down (ignore | red | yellow);
    linkdown (ignore | red | yellow);
    pic-hold-reset (ignore | red | yellow);
    pic-reset (ignore | red | yellow);
    rx-errors (ignore | red | yellow);
    sw-down (ignore | red | yellow);
    tx-errors (ignore | red | yellow);
}
sonet {
    (ais-l | ais-p | ber-sd | ber-sf | locd | lol | lop-p | los | pll | plm-p | rfi-l | rfl-p | uneq-p)
    (ignore | red | yellow);
}
t3 {
    (ais | exz | ferf | idle | lcv | lof | los | pll | ylw) (ignore | red | yellow);
}
}
cluster {
    control-link-recovery;
    control-ports {
        fpc slot-number port port-number;
    }
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    redundancy-group {
        ... the redundancy-group subhierarchy appears at the end of the [edit chassis cluster]
        hierarchy ...
    }
}
reth-count number;
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
}

```

```

    flag flag;
    level severity;
    no-remote-trace;
}
redundancy-group group-number {
    gratuitous-arp-count number;
    hold-down-interval seconds;
    interface-monitor {
        interface-name weight number;
    }
    ip-monitoring {
        family {
            inet {
                ipv4-address {
                    interface rethindex.logical-unit-number secondary-ip-address ipv4-address;
                    weight number;
                }
            }
        }
        global-threshold number;
        global-weight number;
        retry-count count;
        retry-interval interval;
    }
    node node-number priority priority-number;
    preempt;
}
config-button {
    no-clear;
    no-rescue;
}
container-devices {
    device-count number;
}
craft-lockout;
disable-power-management;
disk-partition partition-name (/config | /var) {
    level (full | high) {
        free-space threshold-value (mb | percent);
    }
}
enhanced-policer;
extended-statistics;
fabric {
    degraded {
        action-fpc-restart-disable;
        degraded-fabric-detection-enable
        degraded-fpc-bad-plane-threshold number-bad-planes;
    }
    redundancy-mode (increased-bandwidth | redundant);
}
filter;
fpc slot-number {
    ... the fpc subhierarchy appears after the main [edit chassis] hierarchy ...
}
fpc-feb-connectivity {

```

```
fpc slot-number feb (slot-number | none);
}
fpc-resync;
fru-poweron-sequence sequence;
lcc index {
    ... the lcc subhierarchy appears after the main [edit chassis] hierarchy ...
}
maximum-ecmp value;
memory-enhanced {
    filter;
    route;
    vpn-label;
}
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip);
(packet-scheduling | no-packet-scheduling);
pem {
    minimum number;
}
policer-drop-probability-low;
ppp-subscriber-services (disable | enable);
redundancy {
    cfeb slot (always | preferred);
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    feb {
        redundancy-group group-name {
            description description;
            feb slot-number <backup | primary>;
            no-auto-failover;
        }
    }
    graceful-switchover;
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
route-memory-enhanced;
route-localization {
    inet (chassis);
    inet6;
}
routing-engine {
    bios {
        no-auto-upgrade;
    }
    on-disk-failure disk-failure-action (halt | reboot);
}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
```

```

(source-route | no-source-route);
state [
    cb-upgrade [on | off];
}
synchronization { # for M Series and T Series routers
    primary (external-a | external-b);
    secondary (external-a | external-b);
    signal-type (e1 | t1);
    switching-mode (non-revertive | revertive);
    transmitter-enable;
    validation-interval seconds;
    y-cable-line-termination;
}
synchronization { # for MX80 and MX240 routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality|received-quality);
    source {
        (external-a | external-b) {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
        }
        interfaces interface-name {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
            wait-to-restore minutes;
        }
    }
    switchover-mode (revertive | non-revertive);
}
synchronization { # for ACX Series routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
        (bits | gps) {

```

```

    priority number;
    quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
    request (force-switch | lockout);
  }
  interfaces interface-name {
    priority number;
    quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
    request (force-switch | lockout);
    wait-to-restore minutes;
  }
}
switchover-mode(non-revertive | revertive);
}
system-domains {
  protected-system-domains psdnumerical-index {
    control-plane-bandwidth-percent percent;
    control-slot-numbers [ slot-numbers ];
    control-system-id control-system-id;
    description description;
    fpcs [ slot-numbers ];
  }
  root-domain-id root-domain-id;
}
vrf-mtu-check;
}

chassis {
  fpc slot-number {
    number-of-ports active-ports;
    offline;
    pic slot-number {
      ... the pic subhierarchy appears after the main [edit chassis fpc slot-number] hierarchy
      ...
    }
    port-mirror-instance port-mirror-instance-name;
    power (off | on);
    sampling-instance instance-name;
  }

  fpc slot-number {
    pic slot-number {
      adaptive-services {
        service-package (layer-2 | layer-3 | ...the following extension-provider subhierarchy
          ...);
        extension-provider {
          control-cores number;
          data-cores number;
          data-flow-affinity {
            hash-key (layer-3 | layer-4);
          }
          channelization;
          forwarding-db-size megabytes;
          object-cache-size megabytes;
          package package-name;
          policy-db-size megabytes;
          syslog {

```



```

        facility {
            severity;
            destination (pic-console | routing-engine);
        }
    }
    wired-process-mem-size megabytes;
}
aggregated-devices {
    ima {
        device-count number;
    }
}
aggregate-ports;
atm-cell-relay-accumulation;
atm-l2circuit-mode (aal5 | cell | trunk trunk);
cel {
    e1 port-number {
        channel-group group-number timeslots slot-number;
    }
}
ct3 {
    port port-number {
        t1 link-number {
            channel-group group-number timeslots slot-number;
        }
    }
}
ethernet {
    pic-mode (enhanced-switching | routing | switching);
}
fibre-channel {
    port port-number;
    port-range port-range-low port-range-high
}
egress-policer-overhead bytes;
forwarding-mode {
    sa-multicast;
    vlan-steering {
        vlan-rule (high-low | odd-even);
    }
}
framing (e1 | e3 | sdh | sonet | t1 | t3);
idle-cell-format {
    itu-t;
    payload-pattern payload-pattern-byte;
}
ingress-policer-overhead bytes;
inline-services {
    bandwidth (1g | 10g);
}
linerate-mode;
max-queues-per-interface (4 | 8);
mlfr-uni-nni-bundles number;
no-concatenate;
no-multi-rate;

```

```

port port-number {
    framing (e1 | e3 | sdh | sonet | t1 | t3);
    forwarding-mode {
        sa-multicast;
    }
    speed ( oc3-stm1 | oc12-stm4 | oc48-stm16);
}
port-mirror-instance port-mirror-instance-name;
q-pic-large-buffer {
    (large-scale | small-scale);
}
red-buffer-occupancy {
    weighted-averaged <instant-usage-weight-exponent weight-value>;
}
shdsl {
    pic-mode (1-port-atm | 2-port-atm);
}
sparse-dlcis;
traffic-manager {
    egress-shaping-overhead number;
    ingress-shaping-overhead number;
    mode {
        egress-only;
        ingress-and-egress;
        session-shaping;
    }
}
tunnel-queuing;
tunnel-services {
    bandwidth (1g | 10g | 20g | 40g);
    tunnel-only;
}
vtmapping (itu-t | klm);
}
}

chassis {
    lcc index {
        fpc slot-number {
            ... the fpc subhierarchy appears after the main [edit chassis lcc index] hierarchy ...
        }
        offline;
        online-expected;
    }
}

lcc index {
    fpc slot-number {
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis lcc index fpc slot-number] hierarchy ...
        }
        power (off | on);
        sampling-instance instance-name;
    }
}

```

```

fpc slot-number {
  pic slot-number {
    aggregate-ports;
    atm-cell-relay-accumulation;
    atm-l2circuit-mode (aal5 | cell | trunk trunk);
    framing (e1 | e3 | sdh | sonet | t1 | t3);
    idle-cell-format {
      itu-t;
      payload-pattern payload-pattern-byte;
    }
    linerate-mode;
    max-queues-per-interface (4 | 8);
    no-concatenate;
    no-mcast-replication;
    no-pre-classifier;
    port port-number {
      framing (e1 | e3 | sdh | sonet | t1 | t3);
    }
    q-pic-large-buffer {
      (large-scale | small-scale);
    }
    red-buffer-occupancy {
      weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
      pic-mode (1-port-atm | 2-port-atm);
    }
    traffic-manager {
      egress-shaping-overhead bytes;
      ingress-shaping-overhead bytes;
      mode {
        egress-only;
        ingress-and-egress;
      }
    }
  }
}

```

Related Documentation • Notational Conventions Used in Junos OS Configuration Hierarchies

[\[edit dynamic-profiles\] 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 forwarding-options rpf-loose-mode-discard] Hierarchy Level

```
rpf-loose-mode-discard {
  family {
    inet;
    inet6;
  }
}
```

**Related
Documentation**

- Notational Conventions Used in Junos OS Configuration Hierarchies
- [edit forwarding-options] 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);
(loopback | 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 mlfrr-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

[\[edit multi-chassis\] Hierarchy Level](#)

```

multi-chassis {
  multi-chassis-protection ipv4-address {
    interface interface-name;
  }
}

```

**Related
Documentation**

- Notational Conventions Used in Junos OS Configuration Hierarchies

[\[edit protocols isis\] Hierarchy Level](#)

The following statement hierarchy can also be included at the [\[edit protocols isis\]](#) hierarchy level.

```

protocols {
  isis {

```

```
disable;
clns-routing;
context-identifier ip-address</prefix> {
    level (1 | 2) <disable>;
}
export [ policy-names ];
graceful-restart {
    disable;
    helper-disable;
    restart-duration seconds;
}
ignore-attached-bit;
interface interface-name {
    ... the interface subhierarchy appears after the main [edit protocols isis] hierarchy ...
}
label-switched-path name level level metric metric;
level (1 | 2) {
    disable;
    authentication-key key;
    authentication-type authentication;
    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    prefix-export-limit number;
    wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
no-ipv6-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
    inet6 group-name;
}
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
topologies {
    ipv4-multicast;
    ipv6-multicast;
    ipv6-unicast;
}
traceoptions {
```



```

    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  traffic-engineering {
    disable;
    family inet {
      shortcuts {
        multicast-rpf-routes;
      }
    }
    family inet6 {
      shortcuts;
    }
  }
  ignore-lsp-metrics;
}

isis {
  interface interface-name {
    disable;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    checksum;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | loose | strict);
    ldp-synchronization {
      disable;
      hold-time seconds;
    }
    level (1 | 2) {
      disable;
      hello-authentication-key key;
      hello-authentication-type authentication;
      hello-interval seconds;
      hold-time seconds;
      ipv4-multicast-metric number;
      ipv6-multicast-metric number;
    }
  }
}

```

```
        ipv6-unicast-metric number;  
        metric metric;  
        passive;  
        priority number;  
        te-metric metric;  
    }  
    link-protection;  
    lsp-interval milliseconds;  
    mesh-group (value | blocked);  
    no-adjacency-down-notification;  
    no-eligible-backup;  
    no-ipv4-multicast;  
    no-ipv6-multicast;  
    no-ipv6-unicast;  
    no-unicast-topology;  
    node-link-protection;  
    passive;  
    point-to-point;  
} } }
```

**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 {
  ... 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 15](#)
- Configuring a Layer 2 Control Protocol Routing Instance


bandwidth (Interfaces)

Syntax	<code>bandwidth rate;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure an informational-only bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.
<div>  <p>NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the <code>bandwidth</code> statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$ <p>where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the <code>bandwidth</code> statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.</p> </div>	
Options	<p>rate—Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.</p> <p>Range: Not limited.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Interface Bandwidth

chassis-id


Syntax	<code>chassis-id <i>chassis-id</i>;</code>
Hierarchy Level	[edit]interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the chassis ID of the multichassis aggregated Ethernet interface device. LACP uses the chassis ID to calculate the port number of the multichassis link aggregation group (MC-LAG) physical member links.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

filter

Syntax	<pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<div>  <p>NOTE: On EX Series switches, the <code>group</code>, <code>input-list</code>, <code>output-filter</code> statements are not supported under the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet6</i>], and [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>mpls</i>] hierarchies.</p> </div> <p>Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the family ethernet-switching, inet, inet6, mpls, or vpls only.</p>
Options	<p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. The default filter group number is 0.</p> <p>Range: 0 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Applying a Filter to an Interface Junos Services Interfaces Configuration Release 11.2 Routing Policy Configuration Guide Junos OS System Basics Configuration Guide

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring Firewall Filters (CLI Procedure)
- family

flow-control

Syntax	(flow-control no-flow-control);
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigether-options], [edit interfaces <i>interface-name</i> multiservice-options], [edit interfaces interface-range <i>name</i> aggregated-ether-options], [edit interfaces interface-range <i>name</i> ether-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.
	<div><p>NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.</p></div>
Default	Flow control is enabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Control• Configuring Gigabit Ethernet Interfaces (CLI Procedure)

forward-and-send-to-re

Syntax	forward-and-send-to-re;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface and the Routing Engine. The packets are broadcast only if the egress interface is a LAN interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Targeted Broadcast on page 43 • targeted-broadcast on page 99 • Understanding Targeted Broadcast on page 7

forward-only

Syntax	forward-only;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Targeted Broadcast on page 43 • targeted-broadcast on page 99 • Understanding Targeted Broadcast on page 7

gratuitous-arp-reply

Syntax	(gratuitous-arp-reply no-gratuitous-arp-reply);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Ethernet interfaces, enable updating of the ARP cache for replies received in response to gratuitous ARP requests.
Default	Updating of the ARP cache is disabled on all Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring Gratuitous ARPno-gratuitous-arp-request on page 89


group (RPF Selection)

Syntax	<pre>group group-address{ source source-address{ next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the PIM group address for which you configure RPF selection group (RPF Selection) .
Default	By default, PIM RPF selection is not configured.
Options	group-address —PIM group address for which you configure RPF selection.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring PIM RPF Selection

l2-domain-id-for-l3

Syntax	l2-domain-id-for-l3 <i>id</i> ;
Hierarchy Level	[edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2.
Description	Specify a Layer 2 domain ID within a routing instance.
Options	id —Layer 2 identification number.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring a Layer 2 Virtual Switch on page 15

lACP (Aggregated Ethernet)

Syntax	<pre>lACP { (active passive); admin-key <i>key</i>; fast-failover; link-protection { disable; (revertive non-revertive); } periodic <i>interval</i>; system-id <i>mac-address</i>; system-priority <i>priority</i>; }</pre>
Hierarchy Level	[edit interfaces <i>aex</i> aggregated-ether-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>fast-failover option introduced in Junos OS Release 12.2.</p>
Description	For aggregated Ethernet interfaces only, configure Link Aggregation Control Protocol (LACP).
Default	If you do not specify LACP as either active or <i>passive</i> , LACP remains passive.
Options	<p>active—Initiate transmission of LACP packets.</p> <p>admin-key <i>number</i>—Specify an administrative key for the router or switch.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: You must also configure Multichassis Link Aggregation (MC-LAG) when you configure the admin-key.</p> </div> <p>passive—Respond to LACP packets.</p> <p>fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Aggregated Ethernet LACP Configuring Aggregated Ethernet LACP (CLI Procedure)

- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

layer3-domain-identifier

Syntax	<code>layer3-domain-identifier <i>identifier</i>;</code>
Hierarchy Level	[edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2.
Description	Specify a Layer 3 domain ID within a routing instance.
Options	<code>id</code> —Layer 3 identification number.
Required Privilege Level	<code>system</code> —To view this statement in the configuration. <code>system-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Layer 2 Virtual Switch on page 15

link-protection

Syntax	<pre>link-protection { disable; (revertive non-revertive); }</pre>
Hierarchy Level	<pre>[edit interfaces aex aggregated-ether-options] [edit interfaces aex aggregated-ether-options lacp]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for disable, revertive, and non-revertive statements added in Junos OS Release 9.3.</p>
Description	<p>On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the primary and backup statements at the [edit interfaces ge-fpc/pic/port gigether-options 802.3ad aex] hierarchy level or the [edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex] hierarchy level.</p> <p>On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.</p> <p>For Junos OS link protection, specify link-protection at the following hierarchy levels:</p> <ul style="list-style-type: none">• [edit interfaces ge-fpc/pic/port ether-options 802.3ad aex]• [edit interfaces xe-fpc/pic/port ether-options 802.3ad aex] <p>For LACP standards link protection, specify link-protection at the following hierarchy levels:</p> <ul style="list-style-type: none">• For global LACP link protection, specify at [edit chassis aggregated-devices ethernet lacp]• For a specific aggregated Ethernet interface, specify at [edit interfaces aeX aggregated-ether-options lacp]
Options	The statements are explained separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Aggregated Ethernet Link Protection• Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)

mc-ae-id

Syntax	<code>mc-ae-id <i>mc-ae-id</i>;</code>
Hierarchy Level	[edit interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID.
Options	Range: 1 through 65535.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mode (QFX Series)

Syntax	<code>mode active-active ;</code>
Hierarchy Level	[edit interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the multichassis link aggregation group (MC-LAG) to be in active-active mode. In active-active mode, all of the members of the MC-LAG will be active on both routing or switching devices. Only active-active mode is supported at this time.
Options	active-active —Specify that all of the members of the MC-LAG will be active on both routing or switching devices.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	

mode (Interfaces)

Syntax	mode loose;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Unicast RPF Strict Mode on page 45

multicast-rpf-routes

Syntax	multicast-rpf-routes;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering family inet shortcuts], [edit logical-systems <i>logical-system-name</i> routing-instances traffic-engineering family inet shortcuts], [edit protocols isis traffic-engineering family inet shortcuts], [edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering family inet shortcuts]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Install unicast IPv4 routes into the multicast routing table (inet.2) for multicast reverse-path-forwarding (RPF) checks. Traffic engineering shortcuts must be enabled. IPv4 multicast topology must not be enabled. Label-switched paths (LSPs) must not be advertised into IS-IS.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Enabling IS-IS Traffic Engineering Support

next-hop (PIM RPF Selection)

Syntax	<code>next-hop <i>next-hop-address</i>;</code>
Hierarchy Level	<p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</p>
Release Information	<p>Statement introduced in JUNOS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the specific next-hop address for the PIM group source.
Options	<i>next-hop-address</i> —Specific next-hop address for the PIM group source.
Required Privilege Level	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring PIM RPF Selection



no-gratuitous-arp-request

Syntax	<code>no-gratuitous-arp-request;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p>
Description	For Ethernet interfaces, do not respond to gratuitous ARP requests.
Default	Gratuitous ARP responses are enabled on all Ethernet interfaces.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Gratuitous ARP gratuitous-arp-reply on page 82

no-local-switching

Syntax	no-local-switching;
Hierarchy Level	[edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2.
Description	Specify that access ports in this routing instance do not forward packets to each other.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Layer 2 Virtual Switch on page 15

policer (MAC)

Syntax	<pre> policer { input <i>cos-policer-name</i>; output <i>cos-policer-name</i>; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> accept-source-mac mac-address <i>mac-address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> accept-source-mac mac-address <i>mac-address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), configure MAC policing.</p>
<div>  <p>NOTE: What pics are supported for EX9200?</p> </div>	
<div>  <p>NOTE:</p> <p>On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:</p> <ul style="list-style-type: none"> • Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames. • In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied. • In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule. </div>	
Options	<p>input <i>cos-policer-name</i>—Name of one policer to specify the premium bandwidth and aggregate bandwidth.</p> <p>output <i>cos-policer-name</i>—Name of one policer to specify the premium bandwidth and aggregate bandwidth.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring MAC Address Filtering

prefix-list (PIM RPF Selection)

Syntax	<pre>prefix-list <i>prefix-list-addresses</i> { source <i>source-address</i> { next-hop <i>next-hop-address</i>; } wildcard-source { next-hop <i>next-hop-address</i>; } }</pre>
Hierarchy Level	<pre>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</pre>
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.
Options	<p><i>prefix-list-addresses</i>—List of prefixes (addresses) for multiple PIM groups.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring PIM RPF Selection

redundancy-group

Syntax	<code>redundancy-group <i>group-id</i>;</code>
Hierarchy Level	<code>[edit interfaces aggregated-ether-options mc-ae]</code>
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Multichassis Link Aggregation on page 35

rpf-check (Dynamic Profiles)

Syntax	<pre>rpf-check { mode loose; }</pre>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet protocol family only.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Unicast RPF Strict Mode on page 45

rpf-check (interfaces)

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet or inet6 protocol family only.</p> <p>The mode statement is explained separately.</p>
Options	fail-filter —A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Unicast RPF Strict Mode on page 45• Configuring Unicast RPF Loose Mode on page 46• Example: Configuring Unicast Reverse-Path-Forwarding Check

rpf-check-policy (Routing Options RPF)

Syntax	<code>rpf-check-policy [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.
Options	<i>policy-names</i> —Name of one or more multicast RPF check policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring RPF Policies

rpf-loose-mode-discard

Syntax	<pre>rpf-loose-mode-discard { family { inet; inet6; } }</pre>
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard next hop.
Options	inet —IPv4 address family. inet6 —IPv6 address family.
Required Privilege Level	interface-control—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Unicast RPF on page 44

rpf-selection

Syntax	<pre> rpf-selection { group group-address { source source-address { next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } } prefix-list prefix-list-addresses { source source-address { next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } } } </pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim]
Release Information	<p>Statement introduced in JUNOS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.</p> <p>The remaining statements are explained separately.</p>
Default	If you omit the rpf-selection statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.
Options	source-address —Specific source address for the PIM group.
Required Privilege Level	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Example: Configuring PIM RPF Selection

source (PIM RPF Selection)

Syntax	<code>source source-address { next-hop next-hop-address; }</code>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i>]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the source address for the PIM group.
Options	source-address —Specific source address for the PIM group. The remaining statements are explained separately.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Example: Configuring PIM RPF Selection

status-control

Syntax	<code>status-control (active standby);</code>
Hierarchy Level	[edit interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify whether a peer hosting a multichassis link aggregation group (MC-LAG) is primary or secondary. Primary is considered active, and secondary is considered standby.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

targeted-broadcast

Syntax	targeted-broadcast { forward-and-send-to-re; forward-only; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface. The statements are explained separately.
Default	When this statement is not included, broadcast packets are sent to the Routing Engine only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Targeted Broadcast on page 43 • Understanding Targeted Broadcast on page 7

traceoptions (Individual Interfaces)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>name</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; match; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p>
Description	<p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>The interfaces traceoptions statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system syslog file in the directory /var/log.</p>
Default	If you do not include this statement, no interface-specific tracing operations are performed.
Options	<p>file name—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file files number—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.dcd.</p> <p>match—(Optional) Regular expression for lines to be traced.</p> <p>no-world-readable—(Optional) Prevent any user from reading the log file.</p> <p>world-readable—(Optional) Allow any user to read the log file.</p> <p>size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none">• all—All interface tracing operations• event—Interface events

- **ipc**—Interface interprocess communication (IPC) messages
- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing Operations of an Individual Router or Switch Interface


traps

Syntax	(traps no-traps);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-range name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling or Disabling SNMP Notifications on Physical Interfaces • Enabling or Disabling SNMP Notifications on Logical Interfaces

unicast-reverse-path

Syntax	unicast-reverse-path (active-paths feasible-paths);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-instances <i>routing-instance-name</i> instance-type <i>name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced before Junos OS Release 7.4. Support for routing instances added in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Control the operation of unicast reverse-path-forwarding check. This statement enables the RPF check to be used when routing is asymmetrical.
Options	active-paths —Consider only active paths during the unicast reverse-path check. feasible-paths —Consider all feasible paths during the unicast reverse-path check. Default: If you omit the unicast-reverse-path statement, only the active paths to a particular destination are considered.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Unicast Reverse-Path-Forwarding Check• Enabling Unicast Reverse-Path Forwarding Check for VPNs

unidirectional

Syntax	unidirectional;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Create two new, unidirectional (transmit-only and receive-only) physical interfaces subordinate to the original parent interface. Unidirectional links are currently supported only on 10-Gigabit Ethernet interfaces on the following hardware:
	<div>  <p>NOTE: Which interfaces and pic for EX9200?</p> </div> <ul style="list-style-type: none"> • 4-port 10-Gigabit Ethernet DPC on the MX960 router • 10-Gigabit Ethernet IQ2 PIC and 10-Gigabit Ethernet IQ2E PIC on the T Series router
Default	Disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Unidirectional Traffic Flow on Physical Interfaces • Enabling Unidirectional Traffic Flow on Physical Interfaces

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch• Example: Configuring BGP Autodiscovery for LDP VPLS• Configuring a Layer 3 Subinterface (CLI Procedure)• Configuring Tagged Aggregated Ethernet Interfaces on page 32• Configuring Interfaces for VPLS Routing• Enabling VLAN Tagging• 802.1Q VLANs Overview• vlan-id

wildcard-source (PIM RPF Selection)

Syntax	wildcard-source { next-hop next-hop-address; }
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i>]
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source. The remaining statements are explained separately.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring PIM RPF Selection

PART 3

Administration

- [Routine Monitoring on page 109](#)
- [Operational Commands on page 111](#)

CHAPTER 6

Routine Monitoring

- [Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 109](#)

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface

Purpose To monitor statistics for a Fast Ethernet or Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action `user@host> monitor interface (fe-fpc/pic/port | ge-fpc/pic/port)`



.....

CAUTION: We recommend that you use the `monitor interface fe-fpc/pic/port` or `monitor interface ge-fpc/pic/port` command only for diagnostic purposes. Do not leave these commands on during normal router or switch operations because real-time monitoring of traffic consumes additional CPU and memory resources.

.....

Sample Output

The following sample output is for a Fast Ethernet interface:

```

user@host> monitor interface fe-2/1/0
Interface: fe-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:
  Input bytes:          282556864218 (14208 bps)      [40815]
  Output bytes:         42320313078 (384 bps)         [890]
  Input packets:        739373897 (11 pps)           [145]
  Output packets:       124798688 (1 pps)            [14]
Error statistics:
  Input errors:          0                          [0]
  Input drops:           0                          [0]
  Input framing errors:  0                          [0]
  Policed discards:     6625892                     [6]
  L3 incompletes:       75                          [0]
  L2 channel errors:    0                          [0]
  L2 mismatch timeouts: 0                          [0]
  Carrier transitions:   1                          [0]
  Output errors:         0                          [0]
  Output drops:         0                          [0]
  Aged packets:         0                          [0]
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets       464751787                   [154]
  Packet error count    0                          [0]

```

Meaning Use the information from this command to help narrow down possible causes of an interface problem.



NOTE: If you are accessing the router or switch from the console connection, make sure you set the CLI terminal type using the `set cli terminal` command.

The statistics in the second column are the cumulative statistics since the last time they were cleared using the `clear interfaces statistics interface-name` command. The statistics in the third column are the cumulative statistics since the `monitor interface interface-name` command was executed.

If the input errors are increasing, verify the following:

1. Check the cabling to the router and have the carrier verify the integrity of the line. To verify the integrity of the cabling, make sure that you have the correct cables for the interface port. Make sure you have single-mode fiber cable for a single-mode interface and multimode fiber cable for a multimode interface.
2. For a fiber-optic connection, measure the received light level at the receiver end and make sure that it is within the receiver specification of the Ethernet interface. See Fiber-Optic Ethernet Interface Specifications for the fiber-optic Ethernet interface specifications.
3. Measure the transmit light level on the Tx port to verify that it is within specification. See Fiber-Optic Ethernet Interface Specifications for the optical specifications.

CHAPTER 7

Operational Commands

- [Common Output Fields Description on page 111](#)

Common Output Fields Description

This chapter explains the content of the output fields, which appear in the output of most **show interfaces** commands.

Destination Class Field

For the logical interface, the **Destination class** field provides the names of destination class usage (DCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
	(889)	(597762)
bronze	0	0
	(0)	(0)
silver	0	0
	(0)	(0)

Enabled Field

For the physical interface, the **Enabled** field provides information about the state of the interface, displaying one or more of the following values:

- **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled. To change the interface state to **Enabled**, use the following command:

```
user@host# set interfaces interface enable
```

Manually verify the connections to bring the physical link up.

- **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. To change the interface state to **Enabled**, use the following command:

```
user@host# set interfaces interface enable
```

- **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. Manually verify the connections to bring the physical link up.
- **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.

Filters Field

For the logical interface, the **Filters** field provides the name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is **Filters: Input: *filter-name* and Filters: Output: *filter-name***. For example:

Filters: Input: sample-all
Filters: Output: cp-ftp

Flags Fields

The following sections provide information about flags that are specific to interfaces:

- [Addresses, Flags Field on page 112](#)
- [Device Flags Field on page 113](#)
- [Family Flags Field on page 113](#)
- [Interface Flags Field on page 114](#)
- [Link Flags Field on page 115](#)
- [Logical Interface Flags Field on page 115](#)

Addresses, Flags Field

The **Addresses, Flags** field provides information about the addresses configured for the protocol family on the logical interface and displays one or more of the following values:

- **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status
- **Is-Default**—The default address of the router used as the source address by SNMP, ping, traceroute, and other network utilities.
- **Is-Preferred**—The default local address for packets originating from the local router and sent to destinations on the subnet.
- **Is-Primary**—The default local address for broadcast and multicast packets originated locally and sent out the interface.
- **Preferred**—This address is a candidate to become the preferred address.
- **Primary**—This address is a candidate to become the primary address.

Device Flags Field

The **Device flags** field provides information about the physical device and displays one or more of the following values:

- **Down**—Device has been administratively disabled.
- **Hear-Own-Xmit**—Device receives its own transmissions.
- **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint.
- **Loopback**—Device is in physical loopback.
- **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback.
- **No-Carrier**—On media that support carrier recognition, no carrier is currently detected.
- **No-Multicast**—Device does not support multicast traffic.
- **Present**—Device is physically present and recognized.
- **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.
- **Quench**—Transmission on the device is quenched because the output buffer is overflowing.
- **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering.
- **Running**—Device is active and enabled.

Family Flags Field

The **Family flags** field provides information about the protocol family on the logical interface and displays one or more of the following values:

- **DCU**—Destination class usage is enabled.
- **Dest-route-down**—The software detected that the link is down and has stopped forwarding the link's interface routes.
- **Down**—Protocol is inactive.
- **Is-Primary**—Interface is the primary one for the protocol.
- **Mac-Validate-Loose**—Interface is enabled with loose MAC address validation.
- **Mac-Validate-Strict**—Interface is enabled with strict MAC address validation.
- **Maximum labels**—Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.
- **MTU-Protocol-Adjusted**—The effective MTU is not the configured value in the software.
- **No-Redirects**—Protocol redirects are disabled.

- **Primary**—Interface can be considered for selection as the primary family address.
- **Protocol-Down**—Protocol failed to negotiate correctly.
- **SCU-in**—Interface is configured for source class usage input.
- **SCU-out**—Interface is configured for source class usage output.
- **send-bcast-packet-to-re**—Interface is configured to forward IPv4 broadcast packets to the Routing Engine.
- **targeted-broadcast**—Interface is configured to forward IPv4 broadcast packets to the LAN interface and the Routing Engine.
- **Unnumbered**—Protocol family is configured for unnumbered Ethernet. An unnumbered Ethernet interface borrows an IPv4 address from another interface, which is referred to as the donor interface.
- **Up**—Protocol is configured and operational.
- **uRPF**—Unicast Reverse Path Forwarding is enabled.

Interface Flags Field

The **Interface flags** field provides information about the physical interface and displays one or more of the following values:

- **Admin-Test**—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Hardware-Down**—Interface is nonfunctional or incorrectly connected.
- **Link-Layer-Down**—Interface keepalives have indicated that the link is incomplete.
- **No-Multicast**—Interface does not support multicast traffic.
- **No-receive No-transmit**—Passive monitor mode is configured on the interface.
- **Point-To-Point**—Interface is point-to-point.
- **Pop all MPLS labels from packets of depth**—MPLS labels are removed as packets arrive on an interface that has the **pop-all-labels** statement configured. The depth value can be one of the following:
 - **1**—Takes effect for incoming packets with one label only.
 - **2**—Takes effect for incoming packets with two labels only.
 - **[1 2]**—Takes effect for incoming packets with either one or two labels.
- **Promiscuous**—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
- **Recv-All-Multicasts**—Interface is in multicast promiscuous mode and provides no multicast filtering.

- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Link Flags Field

The **Link flags** field provides information about the physical link and displays one or more of the following values:

- **ACFC**—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.
- **Give-Up**—Link protocol does not continue connection attempts after repeated failures.
- **Loose-LCP**—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.
- **Loose-LMI**—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.
- **Loose-NCP**—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.
- **No-Keepalives**—Link protocol keepalives are disabled.
- **PFC**—Protocol field compression is configured. The PPP session negotiates the PFC option.

Logical Interface Flags Field

The **Logical interface flags** field provides information about the logical interface and displays one or more of the following values:

- **ACFC Encapsulation**—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).
- **Device-down**—Device has been administratively disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Clear-DF-Bit**—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.
- **Hardware-Down**—Interface protocol initialization failed to complete successfully.
- **PFC**—Protocol field compression is enabled for the PPP session.
- **Point-To-Point**—Interface is point-to-point.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Label-Switched Interface Traffic Statistics Field

When you use the **vrf-table-label** statement to configure a VRF routing table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table.

Any routes present in a VRF routing table and configured with the **vrf-table-label** statement are advertised with the LSI logical interface label allocated for the VRF routing table. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table. For more information on the **vrf-table-label** statement, including a list of supported interfaces, see the *Junos VPNs Configuration Guide*.

If you configure the **family mpls** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level and you also configure the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy level, the output for the **show interface interface-name extensive** command includes the following output fields about the LSI traffic statistics:

- **Input bytes**—Number of bytes entering the LSI and the current throughput rate in bits per second (bps).
- **Input packets**—Number of packets entering the LSI and the current throughput rate in packets per second (pps).



NOTE: If LSI interfaces are used with VPLS when **no-tunnel-services** is configured or L3VPN when **vrf-table-label** configuration is applied inside the routing-instance, the **Input packets** field associated with the core-facing interfaces may not display the correct value. Only the Input counter is affected because the LSI is used to receive traffic from the remote PEs. Traffic that arrives on an LSI interface might not be counted at both the Traffic Statistics and the Label-switched interface (LSI) traffic statistics levels.

This note applies to the following platforms:

- M Series routers with -E3 FPC model numbers or configured with an Enhanced CFEB (CFEB-E), and M120 routers
- MX Series routers with DPC or ADPC only

The following example shows the LSI traffic statistics that you might see as part of the output of the **show interface interface-name extensive** command:

```
Label-switched interface (LSI) traffic statistics:
Input bytes:          0          0 bps
Input packets:        0          0 pps
```

Policer Field

For the logical interface, the **Policer** field provides the policers that are to be evaluated when packets are received or transmitted on the interface. The format is **Policer: Input: *type-fpc/picport-in-policer*, Output: *type-fpc/pic/port-out-policer***. For example:

Policer: Input: at-1/2/0-in-policer, Output: at-2/4/0-out-policer

Protocol Field

For the logical interface, the **Protocol** field indicates the protocol family or families that are configured on the interface, displaying one or more of the following values:

- **aenet**—Aggregated Ethernet. Displayed on Fast Ethernet interfaces that are part of an aggregated Ethernet bundle.
- **ccc**—Circuit cross-connect (CCC). Configured on the logical interface of CCC physical interfaces.
- **inet**—IP version 4 (IPv4). Configured on the logical interface for IPv4 protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—IP version 6 (IPv6). Configured on the logical interface for IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), and BGP.
- **iso**—International Organization for Standardization (ISO). Configured on the logical interface for IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). Configured on the logical interface for link services bundling.
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. Configured on the logical interface for multilink bundling.
- **mlppp**—Multilink Point-to-Point Protocol (MLPPP). Configured on the logical interface for multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). Configured on the logical interface for participation in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet (PPPoE). Configured on Ethernet interfaces enabled to support multiple protocol families.
- **tcc**—Translational cross-connect (TCC). Configured on the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol (TNP). Used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only.
- **vpls**—Virtual private LAN service (VPLS). Configured on the logical interface on which you configure VPLS.

RPF Failures Field

For the logical interface, the **RPF Failures** field provides information about the amount of incoming traffic (in packets and bytes) that failed a unicast reverse path forwarding (RPF) check on a particular interface. The format is **RPF Failures: Packets: xx,Bytes: yy**. For example:

RPF Failures: Packets: 0, Bytes:0

Source Class Field

For the logical interface, the **Source class** field provides the names of source class usage (SCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
(889)	(597762)
bronze	0	0
(0)	(0)
silver	0	0
(0)	(0)

clear interfaces statistics

Syntax	clear interfaces statistics (all <i>interface-name</i>)
Release Information	Command introduced before Junos OS Release 7.4.
Description	Set interface statistics to zero. If you issue the clear interfaces statistics <i>interface-name</i> command and then perform a graceful Routing Engine switchover, the interface statistics are not cleared on the new master. Reissue the command to clear the interface statistics again.
Options	all —Set statistics on all interfaces to zero. <i>interface-name</i> —Set statistics on a particular interface to zero.
Required Privilege Level	clear
List of Sample Output	clear interfaces statistics on page 119
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

<code>clear interfaces statistics</code>	<code>user@host> clear interfaces statistics</code>
--	--

show interfaces (10-Gigabit Ethernet)

Syntax	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
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	<p><code><i>xe-fpc/pic/port</i></code>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 135</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 138</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 140</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 143</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 143</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 145</p>
Output Fields	See Table 4 on page 121 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 4: 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” on page 111 .	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: <ul style="list-style-type: none"> • 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” on page 111 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 111 .	All levels

Table 4: 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 ” on page 111.	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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • 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. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 4 on page 121.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • 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)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>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.</p> <ul style="list-style-type: none"> • 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	<p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> • efec—Enhanced forward error correction (EFEC) is configured to detect 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	<p>OTN mode.</p> <ul style="list-style-type: none"> • 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	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> • 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	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect. 	detail extensive
OTN OC	<p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
OTN OTU	OTN OTU defects detected on the interface <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • 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 5 on page 135 • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>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.</p> <ul style="list-style-type: none"> • 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	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • 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	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • 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 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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	<p>(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.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 4: 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. <ul style="list-style-type: none"> • 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" on page 111 .	All levels

Table 4: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> 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:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> 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 ” on page 111.	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	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> 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

Table 4: 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” on page 111 .	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” on page 111 .	detail extensive none
<i>protocol-family</i>	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” on page 111 .	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 4: 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 5 on page 135](#) 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 5 on page 135](#), 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 5: 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
  Bit errors           Seconds
                        0

```



```

    Errored blocks                                0
MAC statistics:
    Receive
    Transmit
    Total octets                                6970332384
    Total packets                                81050506
    Unicast packets                              81050000
    Broadcast packets                             506
    Multicast packets                             0
    CRC/Align errors                             0
    FIFO errors                                  0
    MAC control frames                           0
    MAC pause frames                             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
    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 (Discard)

Syntax	<pre>show interfaces dsc <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified discard interface.
Options	<p>dsc—Display standard information about the specified discard interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—This option is not relevant for the discard interface and always shows a value of 0.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) This option is not relevant for the discard interface and always shows a value of 0.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show interfaces (ATM) • show interfaces routing
List of Sample Output	show interfaces dsc on page 151 show interfaces dsc brief on page 151 show interfaces dsc detail on page 151 show interfaces dsc extensive on page 151
Output Fields	Table 6 on page 147 lists the output fields for the show interfaces (discard) command. Output fields are listed in the approximate order in which they appear.

Table 6: Discard show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, whether the interface is enabled, and the state of the physical interface: Up or Down .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none

Table 6: Discard show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Type of interface. Software-Pseudo indicates a standard software interface with no associated hardware device.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be Internal or External .	brief detail extensive
Speed	Speed at which the interface is running.	brief detail extensive
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 111 .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 111 .	All levels
Link type	Encapsulation being used on the physical interface.	detail extensive
Link flags	Information about the link. Possible values are described in the "Link Flags" section under "Common Output Fields Description" on page 111 .	detail extensive
Physical info	Information about the physical interface.	detail extensive
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive
Current address, Hardware address	Configured MAC address and hardware MAC address.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 6: Discard show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
Output errors	<p>(Extensive only) Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 6: Discard show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 111 .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family configured on the logical interface, such as iso , inet6 , or mpls .	All levels
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive

Sample Output

show interfaces dsc

```
user@host> show interfaces dsc
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5
  Type: Software-Pseudo, MTU: Unlimited
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

Logical interface dsc.0 (Index 66) (SNMP ifIndex 235)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited
  Flags: None
```

show interfaces dsc brief

```
user@host> show interfaces dsc brief
Physical interface: dsc, Enabled, Physical link is Up
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface dsc.0
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  inet
```

show interfaces dsc detail

```
user@host> show interfaces dsc detail
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0

Logical interface dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
  Flags: None
```

show interfaces dsc

```
user@host> show interfaces dsc extensive
Physical interface: dsc, Enabled, Physical link is Up
```

extensive

```
Interface index: 5, SNMP ifIndex: 5, Generation: 9
Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link type      : Unspecified
Link flags     : None
Physical info  : Unspecified
Hold-times     : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :                0
  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 dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
```


show interfaces (Gigabit Ethernet)

Syntax	<pre>show interfaces <i>ge-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
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	<p><i>ge-fpc/pic/port</i>—Display standard information about the specified Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
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	<ul style="list-style-type: none"> Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 169</p> <p>show interfaces (Gigabit Ethernet on MX Series Routers) on page 169</p> <p>show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 170</p> <p>show interfaces brief (Gigabit Ethernet) on page 171</p> <p>show interfaces detail (Gigabit Ethernet) on page 171</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 172</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 175</p> <p>show interfaces (ACI Interface Set Configured) on page 176</p>
Output Fields	<p>Table 7 on page 154 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 8 on page 167.</p>

Table 7: 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” on page 111 .	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: <ul style="list-style-type: none"> • 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” on page 111 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 111 .	All levels

Table 7: 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 ” on page 111.	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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • 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. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>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.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • 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)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • 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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>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.</p> <ul style="list-style-type: none"> • 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	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • 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	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • 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	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • 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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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. <p>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.</p> <ul style="list-style-type: none"> • 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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>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.</p> <ul style="list-style-type: none"> • 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	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • 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	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • 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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • 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. <p>Subfields are:</p> <ul style="list-style-type: none"> • 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 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • 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 " on page 111.	All levels

Table 7: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • 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	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • 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 ” on page 111.	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 7: 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	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • 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	<p>Number and rate of bytes and packets transiting the switch.</p> <p>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.</p>	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" on page 111 .	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 7: 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” on page 111 .	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” on page 111 .	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 8: 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	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	<p>Traffic statistics:</p> <p>Input bytes: 478 bytes per packet, representing the Layer 3 packet</p>	

Table 8: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

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

Sample Output

show interfaces (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues    : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped  : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms : None
  Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress 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:
Total octets          Receive          Transmit
                    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       500000000   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       500000000   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 7 on page 154](#).

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 (Serial)

Syntax	<pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about serial interfaces, including RS-232, RS-422/449, EIA-530, X.21, and V.35.
Options	<p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>se-fpc/pic/port</i>. On the J Series routers, the interface type is <i>se-pim/O/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (Serial, EIA-530) on page 184</p> <p>show interfaces brief (Serial, EIA-530) on page 184</p> <p>show interfaces detail (Serial, EIA-530) on page 184</p> <p>show interfaces extensive (Serial, EIA-530) on page 185</p> <p>show interfaces (Serial, V.35) on page 186</p> <p>show interfaces brief (Serial, V.35) on page 186</p> <p>show interfaces detail (Serial, V.35) on page 187</p> <p>show interfaces extensive (Serial, V.35) on page 187</p> <p>show interfaces statistics detail (RS 449) on page 189</p>
Output Fields	Table 9 on page 177 lists the output fields for the show interfaces (Serial) command. Output fields are listed in the approximate order in which they appear.

Table 9: show interfaces (Serial) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 9: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 111 .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface.	All levels
Maximum speed	Maximum speed. The nonconfigurable value is 16,384 kbps.	detail extensive none
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 111 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 111 .	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 111 .	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Keepalive settings	(PPP and HDLC) Configured settings for keepalive packets. <ul style="list-style-type: none"> Interval <i>seconds</i>—Time between successive keepalive requests. The range of values, in seconds, is 10 to 32,767. The default value is 10. Up-count <i>number</i>—Number of keepalive packets a destination must receive to change a link's status from down to up. The range of values is 1 to 255. The default value is 1. Down-count <i>number</i>—Number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 to 255. The default value is 3. 	All levels
Keepalive	(PPP and HDLC) Information about keepalive packets. <ul style="list-style-type: none"> Input: <i>number (hh:mm:ss ago)</i>—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. Output: <i>number (hh:mm:ss ago)</i>—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent. 	brief none

Table 9: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive statistics	(PPP and HDLC) Information about keepalive packets. <ul style="list-style-type: none"> • Input: <i>number (last seen hh:mm:ssago)</i>—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • Output: <i>number (last seen hh:mm:ss ago)</i>—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent. 	detail extensive
LCP state	(PPP) Link Control Protocol state. <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—LCP negotiation is incomplete (not yet completed or has failed). • Not-configured—LCP is not configured on the interface. • Opened—LCP negotiation is successful. 	detail extensive none
NCP state	(PPP) Network Control Protocol state. <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. 	detail extensive none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. <ul style="list-style-type: none"> • Chap-Chal-received—Challenge was received but response not yet sent. • Chap-Chal-sent—Challenge was sent. • Chap-Resp-received—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • Chap-Resp-sent—Response was sent for the challenge received. • Closed—CHAP authentication is incomplete. • Failure—CHAP authentication failed. • Not-configured—CHAP is not configured on the interface. • Success—CHAP authentication was successful. 	detail extensive none
CoS queues	Number of CoS queues configured.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified

Table 9: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • 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	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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. 	extensive
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeds the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues supported	Total number of egress queues supported on the specified interface. Displayed with the statistics option.	detail extensive
Egress queues in use	Total number of egress queues in use on the specified interface. Displayed with the statistics option.	detail extensive

Table 9: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. Displayed with the statistics option. <ul style="list-style-type: none"> 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
Serial media information	Information about the physical media: <ul style="list-style-type: none"> Line protocol—eia530, eia530a, rs232, rs449, v.35, or x.21. Resync history—Information about resynchronization events: <ul style="list-style-type: none"> Sync loss count—Number of times the synchronization was lost. Data signal—(X.21 and V.35) Information about the data signal: <ul style="list-style-type: none"> Rx Clock—Receive clock status: OK (DTE is receiving the receive clock signal) or Not detected (receive clock signal is not being received). Control signals—Information about modem control signals: <ul style="list-style-type: none"> Local mode:DCE (data communication equipment) or DTE (data terminal equipment) To DCE—Control signals that the Serial PIC sent to the DCE: DTR (Data Terminal Ready:up or down) or RTS (Request To Send: up or down.) From DC—Control signals that the Serial PIC received from the DCE: CTS (Clear To Send: up or down), DCD (Data Carrier Detect: up or down), DSR (Data Set Ready: up or down), or TM (Test Mode: up or down). Clocking mode—Clocking used for the transmit clock: <ul style="list-style-type: none"> dte—Transmit clock is generated by DTE. dce—Transmit clock is generated by the DCE and is looped back as the transmit clock. loop-timed—Receive clock from the DCE is looped back as the transmit clock. Clock rate—Rate, in megahertz (MHz), at which the clock is configured. Loopback—Configured loopback mode for the interface: dce-remote, dce-local, liu, local, or none. Tx clock—Clocking phase of the transmit clock: invert (transmit clock polarity is inverted) or non-invert (transmit clock polarity is not inverted). Line encoding—Type of line encoding used: nrz (nonreturn to zero) or nrzi (return to zero inverted). 	detail extensive
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> Destination slot—FPC slot number. PLP byte—Packet Level Protocol byte. 	extensive

Table 9: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface: <ul style="list-style-type: none"> • 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	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “ Common Output Fields Description ” on page 111.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the source and destination address are also displayed.	brief
Protocol	Protocol family configured on the logical interface, such as iso, inet6, mpls.	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “ Common Output Fields Description ” on page 111.	detail extensive

Table 9: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 111 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 32 (00:00:10 ago), Output: 31 (00:00:07 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues    : 8 supported, 8 maximum usable queues
  Last flapped  : 2006-04-26 15:10:18 PDT (00:05:22 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3
```

show interfaces brief (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 brief
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 235 (00:00:10 ago), Output: 234 (00:00:00 ago)

Logical interface se-5/0/1.0
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  inet 12.0.0.1/30
```

show interfaces detail (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 detail
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Hold-times    : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 37 (last seen 00:00:06 ago)
    Output: 35 (last sent 00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues    : 8 supported, 8 maximum usable queues
  Last flapped  : 2006-04-26 15:10:18 PDT (00:06:02 ago)
```



```

Statistics last cleared: Never
Traffic statistics:
  Input bytes :          928          40 bps
  Output bytes :        1023          48 bps
  Input packets:          76           0 pps
  Output packets:        77           0 pps
Serial media information:
  Line protocol: eia530
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: OK
  Control signals:
    Local mode: DTE
    To DCE: DTR: up, RTS: up
    From DCE: CTS: up, DCD: up, DSR: up
  Clocking mode: loop-timed
  Clock rate: 8.0 MHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 15, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
  Generation: 23

```

show interfaces extensive (Serial, EIA-530)

```

user@host> show interfaces se-5/0/1 extensive
Physical interface: se-5/0/1, Enabled, Physical link is Up
Interface index: 144, SNMP ifIndex: 41, Generation: 25
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
Device flags : Present Running
Interface flags: Point-To-Point Internal: 0x4000
Link flags : Keepalives
Hold-times : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 40 (last seen 00:00:00 ago)
  Output: 37 (last sent 00:00:09 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
CoS queues : 8 supported, 8 maximum usable queues
Last flapped : 2006-04-26 15:10:18 PDT (00:06:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          988          40 bps
  Output bytes :        1088          48 bps
  Input packets:          81           0 pps
  Output packets:        82           0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 2, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Serial media information:

```

```

Line protocol: eia530
Resync history:
  Sync loss count: 0
Data signal:
  Rx Clock: OK
Control signals:
  Local mode: DTE
  To DCE: DTR: up, RTS: up
  From DCE: CTS: up, DCD: up, DSR: up
Clocking mode: loop-timed
Clock rate: 8.0 MHz
Loopback: none
Tx clock: non-invert
Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue      Bandwidth      Buffer  Priority  Limit
                           %      bps      %      usec
0 best-effort             95      15564800  95        0      low     none
3 network-control         5       819200   5         0      low     none

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 15, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
  Generation: 23

```

show interfaces (Serial, V.35)

```

user@host> show interfaces se-5/0/0
Physical interface: se-5/0/0, Enabled, Physical link is Down
Interface index: 150, SNMP ifIndex: 39
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
Link flags     : Loose-NCP
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 0 (never)
LCP state: Down
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Closed
CoS queues     : 8 supported, 8 maximum usable queues
Last flapped   : 2006-04-26 14:51:27 PDT (01:02:23 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Protocol inet, MTU: 1500
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3

```

show interfaces brief (Serial, V.35)

```

user@host> show interfaces se-5/0/0 brief
Physical interface: se-5/0/0, Enabled, Physical link is Down
Type: Serial, Link-level type: PPP, MTU: 1504
Device flags   : Present Running Down

```

```

Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
Link flags      : Loose-NCP
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 0 (never)

```

```

Logical interface se-5/0/0.0
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  inet 13.0.0.2/30

```

show interfaces detail (Serial, V.35)

```

user@host> show interfaces se-5/0/0 detail
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39, Generation: 31
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags     : Loose-NCP
  Hold-times    : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
    Output: 0 (last sent: never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
  CoS queues   : 8 supported, 8 maximum usable queues
  Last flapped : 2006-04-26 14:51:27 PDT (01:03:15 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
  Serial media information:
    Line protocol: v.35
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: Not Detected
    Control signals:
      Local mode: DCE
      To DTE: CTS: down, DCD: down, DSR: up
      From DTE: DTR: down, RTS: down
    DCE loopback override: Off
    Clocking mode: internal
    Clock rate: 38.4 KHz
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  Protocol inet, MTU: 1500, Generation: 17, Route table: 0
  Flags: Protocol-Down
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
    Generation: 23

```

show interfaces extensive (Serial, V.35)

```

user@host> show interfaces se-5/0/0 extensive
Physical interface: se-5/0/0, Enabled, Physical link is Down
Interface index: 150, SNMP ifIndex: 39, Generation: 31
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
Link flags     : Loose-NCP
Hold-times     : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 0 (last seen: never)
  Output: 0 (last sent: never)
LCP state: Down
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Closed
CoS queues   : 8 supported, 8 maximum usable queues
Last flapped : 2006-04-26 14:51:27 PDT (01:04:17 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, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Serial media information:
  Line protocol: v.35
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: Not Detected
  Control signals:
    Local mode: DCE
    To DTE: CTS: down, DCD: down, DSR: up
    From DTE: DTR: down, RTS: down
  DCE loopback override: Off
  Clocking mode: internal
  Clock rate: 38.4 KHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %          bps      %      usec
  0 best-effort           95        15564800   95         0        low   none
  3 network-control        5         819200     5          0        low   none

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 17, Route table: 0
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,

```

Generation: 23

show interfaces statistics detail (RS 449)

```

user@host> show interfaces se-6/0/0 statistics detail
Interface index: 149, SNMP ifIndex: 59, Generation: 150
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 8mbps
Device flags : Present Running
Interface flags: Point-To-Point Internal: 0x4000
Link flags : No-Keepalives Loose-NCP
Hold-times : Up 0 ms, Down 0 ms
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
CoS queues : 8 supported, 8 maximum usable queues
Last flapped : 2007-11-28 19:38:36 PST (00:14:06 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 744 0 bps
Output bytes : 5978 0 bps
Input packets: 33 0 pps
Output packets: 129 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
Resource errors: 0
Output errors:
Carrier transitions: 13, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Egress queues: 8 supported, 5 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 24 24 0
1 expedited-fo 0 0 0
2 bulk 0 0 0
3 assured-forw 105 105 0
4 voip 0 0 0

Serial media information:
Line protocol: rs449
Resync history:
Sync loss count: 0
Data signal:
Rx Clock: OK
Control signals:
Local mode: DTE
To DCE: DTR: up, RTS: up
From DCE: CTS: up, DCD: up, DSR: up
Clocking mode: internal
Loopback: none
Tx clock: non-invert
Line encoding: nrz

Logical interface se-6/0/0.0 (Index 75) (SNMP ifIndex 69) (Generation 141)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 256, Generation: 145, Route table: 0
Flags: None

```

Addresses, Flags: Is-Preferred Is-Primary
Destination: 11.11.11/24, Local: 11.11.11.2, Broadcast: 11.11.11.255,
Generation: 157

show interfaces extensive

Syntax show interfaces extensive

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.

Description Display extensive information about all interfaces configured on the router.



NOTE:

- At some times, the cumulative byte counters displayed with the `show interfaces extensive` command on the 10-Gigabit Ethernet MPC with SFP+ is not always increasing and cumulative and does not give the correct results. There is a time lag in collecting these statistics, during which the display might decrease or go from a nonzero number to zero. Eventually, the counter will display the correct result.
- When the `show interfaces extensive` command is executed on a router with an MPC or a T4000 Type 5 FPC, the *Input packet rejects* counter of the *Filter statistics* field also displays statistics related to the following packet errors:
 - Invalid VLAN range
 - Tagged packet received on an untagged interface
- When the `show interfaces extensive` command is executed on an interface that is configured on a T4000 Type 5 FPC, the *IPv6 transit statistics* field displays:
 - Total statistics (sum of transit and local statistics) at the physical interface level
 - Transit statistics at the logical interface level
- When the `show interfaces extensive` command is executed on an aggregate interface in a T1600 Core Router, the *IPv6 Input bytes* is displayed for an aggregate interface. However, the *IPv6 Input bytes* is always zero on a member link of an aggregated bundle even when there are IPv6 transit traffic on the member link. This is because the logical interface index of the aggregate logical interface is updated but not the logical interface of the member links in the channel lookup table.

Options This command has no options.

Required Privilege Level view

List of Sample Output [show interfaces extensive \(Circuit Emulation\) on page 193](#)

[show interfaces extensive \(Fast Ethernet\) on page 193](#)
[show interfaces extensive \(Gigabit Ethernet\) on page 195](#)
[show interfaces extensive \(10-Gigabit Ethernet\) on page 195](#)
[show interfaces extensive \(IQ2 and IQ2E\) on page 198](#)
[show interfaces extensive \(100-Gigabit Ethernet\) on page 201](#)
[show interfaces extensive \(PTX5000 Packet Transport Switch\) on page 202](#)
[show interfaces extensive \(T4000 Routers with Type 5 FPCs\) on page 204](#)
[show interfaces extensive \(T4000 Routers with 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC\) on page 206](#)
[show interfaces extensive \(Aggregated Ethernet\) on page 208](#)

Output Fields For more information, see the output fields table for the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under [“Common Output Fields Description” on page 111](#). For sample output for specific interfaces, see the other topics in this collection.

Sample Output

show interfaces extensive (Circuit Emulation)

If a Circuit Emulation (CE) PIC is configured for SAToP pseudowire, then pseudowire statistics are displayed in the CE information section of the show interface extensive output. If SAToP pseudowire is not configured on the CE PIC, then all the CE information counters will be displayed as 0 (zero).

```
user@host> show interface t1-0/0/0 extensive
Physical interface :t1-0/0/0, Enabled, Physical Link : Up
  Interface index:61441
  Speed : 1.54 Mbps, Loopback: Disabled
  Operational state : Enabled,   Encapsulation : Trans
  Encoding : b8zs,           Framing : unframe,   Build-out : 0-30
  Inversion : enable, Clock source : master
  Description :
  Traffic statistics:
  T1 media:      Seconds
  ES              1643
  SES            1643

  CE Info      Packets      Bytes
  CE Rx       : 2395529     306627712
  CE Tx       : 2396259     306721152
  CE Rx Drop: 0             0
  CE Tx Drop: 0             0

  CE Overrun  Events: 0
  CE Underrun Events: 0
```

Sample Output

show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/2/1 extensive
Physical interface: fe-0/2/0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:91:c4:3e, Hardware address: 00:90:69:91:c4:3e
  Last flapped   : 2006-04-16 23:00:41 PDT (02:08:05 ago)
  Statistics last cleared: 2006-04-16 21:42:00 PDT (03:26:46 ago)
  Traffic statistics:
  Input bytes   : 17539      152 bps
  Output bytes  : 92968      224 bps
  Input packets: 348         0 pps
  Output packets: 1349       0 pps
  Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
  Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 4 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped packets
```

0 best-effort	66	66	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	1283	1283	0

Active alarms : None

Active defects : None

MAC statistics:

	Receive	Transmit
Total octets	24721	105982
Total packets	348	1349
Unicast packets	347	430
Broadcast packets	1	37
Multicast packets	0	882
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	348	
Input packet rejects	0	
Input DA rejects	0	
Input SA rejects	0	
Output packet count		1349
Output packet pad count		0
Output packet error count		0
CAM destination filters: 3, CAM source filters: 0		

Autonegotiation information:

Negotiation status: Complete

Link partner:

Link mode: Full-duplex, Flow control: None, Remote fault: OK

Packet Forwarding Engine configuration:

Destination slot: 0

CoS information:

CoS transmit queue	Bandwidth	Buffer	Priority	Limit
	%	bps	%	usec
0 best-effort	95	95000000	95	0
3 network-control	5	5000000	5	0
			low	none
			low	none

Logical interface fe-0/2/0.0 (Index 66) (SNMP ifIndex 46) (Generation 133)

Flags: SNMP-Traps Encapsulation: ENET2

Protocol inet, MTU: 1500, Generation: 142, Route table: 0

Flags: DCU, SCU-out

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
silv1_new	0	0
(0)	(
silv2_new	0	0
(0)	(
silv_misc	0	0
(0)	(
silver0	0	0
(0)	(

```

silver2          0          0
(                0) (      0)
silver3          0          0
(                0) (      0)
silver4          0          0
(                0) (      0)
silver5          0          0
(                0) (      0)
silver6          0          0
(                0) (      0)
silver7          0          0
(                0) (      0)
silver9          0          0
(                0) (      0)
Source class      Packets      Bytes
                  (packet-per-second) (bits-per-second)
gold1            0          0
(                0) (      0)
gold2            16600      1062400
(                0) (      0)
gold3            0          0
(                0) (      0)
Addresses, Flags: Is-Preferred Is-Primary
Destination: 12.1.1/24, Local: 12.1.1.1, Broadcast: 12.1.1.255,
Generation: 150

```

Sample Output

show interfaces
extensive (Gigabit
Ethernet)

```
user@host> show interfaces ge-5/0/0.0 extensive
```

```

Logical interface ge-5/0/0.0 (Index 71) (SNMP ifIndex 1930) (Generation 139)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Traffic statistics:
  Input bytes :          0
  Output bytes :         42
  Input packets:          0
  Output packets:         1
Local statistics:
  Input bytes :          0
  Output bytes :         42
  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
Output Filters: f-any
Protocol inet, MTU: 1500, Generation: 155, Route table: 0
  Output Filters: f-inet,
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.1/24, Local: 10.11.1.1, Broadcast: 10.11.1.255,
    Generation: 170
Protocol multiservice, MTU: Unlimited, Generation: 156, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer__

```

Sample Output

**show interfaces
extensive (10-Gigabit
Ethernet)**

```

user@host> show interfaces xe-2/1/0 extensive

Physical interface: xe-2/1/0, Enabled, Physical link is Up
  Interface index: 258, SNMP ifIndex: 762, Generation: 2046
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
  None, Loopback: None, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f8:6d:eb, Hardware address: 00:1d:b5:f8:6d:eb
  Last flapped   : 2011-12-17 00:19:02 PST (07:36:37 ago)
  Statistics last cleared: 2011-12-17 07:55:24 PST (00:00:15 ago)
  Traffic statistics:
    Input bytes :          110000          0 bps
    Output bytes :           0          0 bps
    Input packets:           1000          0 pps
    Output packets:           0          0 pps
  IPv6 transit statistics:
    Input bytes :          110000
    Output bytes :           0
    Input packets:           1000
    Output packets:           0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0,
    MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort          0              0              0
    1 expedited-fo        0              0              0
    2 assured-forw        0              0              0
    3 network-cont        0              0              0

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                expedited-forwarding
    2                assured-forwarding
    3                network-control

  Active alarms : None
  Active defects : None
  PCS statistics
    Bit errors          0
    Errored blocks      0
  MAC statistics:
    Receive              Transmit
    Total octets        128000          0
    Total packets       1000          0
    Unicast packets     1000          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                      1000
  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: 2
CoS information:
  Direction : Output
  CoS transmit queue                    Bandwidth          Buffer Priority
Limit
      0 best-effort                      95      95000000000    95      0      low
none
      3 network-control                   5      5000000000      5      0      low
none
Interface transmit statistics: Disabled

Logical interface xe-2/1/0.0 (Index 83) (SNMP ifIndex 1677) (Generation 10082)

Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input bytes :                          110000
  Output bytes :                          0
  Input packets:                          1000
  Output packets:                          0
IPv6 transit statistics:
  Input bytes :                          55000
  Output bytes :                          0
  Input packets:                          500
  Output packets:                          0
Local statistics:
  Input bytes :                          55000
  Output bytes :                          0
  Input packets:                          500
  Output packets:                          0
Transit statistics:
  Input bytes :                          55000
  Output bytes :                          0
  Input packets:                          500
  Output packets:                          0
IPv6 transit statistics:
  Input bytes :                          55000
  Output bytes :                          0
  Input packets:                          500
  Output packets:                          0
Protocol inet6, MTU: 1500, Generation: 23739, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 2001:1000:abcd:2312:1432:abcd:1234:0/112, Local:

```

```

2001:1000:abcd:2312:1432:abcd:1234:1234
  Generation: 506
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::21d:b5ff:fe8:6deb
  Protocol multiservice, MTU: Unlimited, Generation: 508
  Generation: 23740, Route table: 0
  Policer: Input: __default_arp_policer__

```

Sample Output

show interfaces
extensive (IQ2 and
IQ2E)

```

user@host> show interfaces ge-3/2/2 extensive
Physical interface: ge-3/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 548, Generation: 159
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues    : 8 supported, 8 maximum usable queues
Schedulers    : 128
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:14:f6:12:86:fa, Hardware address: 00:14:f6:12:86:fa
Last flapped   : 2010-03-17 04:03:11 PDT (00:45:30 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :          1716096          0 bps
  Output bytes  :          1716448          0 bps
  Input packets :           13407          0 pps
  Output packets:           13411          0 pps
IPv6 total statistics:
  Input bytes   :          1716096
  Output bytes  :          1716096
  Input packets :           13407
  Output packets:           13407
Ingress traffic statistics at Packet Forwarding Engine:
  Input bytes   :          1716096          0 bps
  Input packets :           13407          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: 1, 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: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped
packets
  0 best-effort          13407          13407
0
  1 expedited-fo              0              0
0
  2 assured-forw              0              0
0
  3 network-cont              0              0
0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped
packets

```

```

0      0 best-effort          13407          13407
0      1 expedited-fo        0              0
0      2 assured-forw        0              0
0      3 network-cont        4              4
0
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive          Transmit
Total packets        1716096        1716448
Unicast packets      13407         13411
Broadcast packets    0              4
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    13407
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   13411
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: None, Remote fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 3
CoS information:
Direction : Output
CoS transmit queue    Bandwidth          Buffer Priority
Limit
0 best-effort          %          bps          %          usec          low
none                    95    950000000    95            0
3 network-control      5      50000000    5            0          low
none
Direction : Input
CoS transmit queue    Bandwidth          Buffer Priority
Limit
0 best-effort          %          bps          %          usec          low
none                    95    950000000    95            0
3 network-control      5      50000000    5            0          low
none

```

Logical interface ge-3/2/2.0 (Index 83) (SNMP ifIndex 6080) (Generation

148)

```

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Traffic statistics:
  Input bytes :          0
  Output bytes :         336
  Input packets:          0
  Output packets:         4
IPv6 total statistics:
  Input bytes :       1716096
  Output bytes :       1716096
  Input packets:      13407
  Output packets:     13407
Local statistics:
  Input bytes :          0
  Output bytes :         336
  Input packets:          0
  Output packets:         4
Transit statistics:
  Input bytes :          0          0 bps
  Output bytes :          0          0 bps
  Input packets:          0          0 pps
  Output packets:          0          0 pps
IPv6 total statistics:
  Input bytes :       1716096
  Output bytes :       1716096
  Input packets:      13407
  Output packets:     13407
Protocol inet6, MTU: 1500, Generation: 159, Route table: 0
  Flags: Is-Primary
  Addresses, Flags: Is-Default Is-Primary
    Destination: Unspecified, Local: 2000::2
Generation: 146
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::214:f600:6412:86fa
Protocol multiservice, MTU: Unlimited, Generation: 148
Generation: 160, Route table: 0
  Policer: Input: __default_arp_policer__

```

Logical interface ge-3/2/2.32767 (Index 84) (SNMP ifIndex 6081) (Generation

149)

```

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
Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```


Sample Output

**show interfaces
extensive (100-Gigabit
Ethernet)**

```

user@host> show interfaces et-0/0/0:0 extensive
Physical interface: et-0/0/0:0, Enabled, Physical link is Down
  Interface index: 156, SNMP ifIndex: 516, Generation: 163
  Link-level type: Ethernet, MTU: 9192, Speed: 50000mbps, 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
  CoS queues    : 8 supported, 8 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:aa:aa:aa:aa:00, Hardware address: 00:21:59:5c:48:00
  Last flapped   : 2010-01-07 16:36:49 PST (18:02:35 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: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0,
  HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 DEFAULT, NC-                0                0                0
  1 REALTIME                    0                0                0
  2 PRIVATE, NC-                0                0                0
  3 CONTROL                     1253               1253                0
  4 BC-H, CLASS_                0                0                0
  5 BC-M, CLASS_                0                0                0
  6 IA, CLASS_V_                0                0                0
  7 CLASS_S_OUTP                0                0                0

Queue      Mapped Forwarding Class
0          DEFAULT, NC-Q0
1          REALTIME
2          PRIVATE, NC-Q1
3          CONTROL
4          BC-H, CLASS-Q4
5          BC-M, CLASS-Q5

```

```

6      IA, CLASS_V_OUTPUT
7      CLASS_S_OUTPUT
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive      Transmit
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
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority Limit

                                %      bps      %      usec
0 best-effort         95    47500000000    95      0      low none
3 network-control     5     25000000000     5      0      low none

Logical interface et-0/0/0:0.0 (Index 68) (SNMP ifIndex 546) (Generation 161)
Flags: Deviet-Down SNMP-Traps 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
Protocol inet, MTU: 9178, Generation: 220, Route table: 0
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 210.160.0/24, Local: 210.160.0.1, Broadcast: 210.160.0.255,
Generation: 192
Protocol mpls, MTU: 9166, Maximum labels: 3, Generation: 221, Route table: 0

Protocol multiservice, MTU: Unlimited, Generation: 222, Route table: 0
Policer: Input: __default_arp_policer

```

Sample Output

show interfaces

```

user@host> show interfaces et-7/0/0 extensive
Physical interface: et-7/0/0, Enabled, Physical link is Up

```

extensive (PTX5000
Packet Transport
Switch)

```

Interface index: 168, SNMP ifIndex: 501, Generation: 171
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
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: 88:e0:f3:3b:de:43, Hardware address: 88:e0:f3:3b:de:43
Last flapped : 2012-01-18 11:48:24 PST (01:47:08 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 3583014 0 bps
Output bytes : 758050 0 bps
Input packets: 17740 0 pps
Output packets: 3418 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, Runt: 0, Policed discards: 0, L3
incompletes: 0,
L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors:
0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 252 252 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 6196 6196 0

Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control
Active alarms : None
Active defects : None
MAC statistics:
Total octets 4108825 1159686
Total packets 21166 6448
Unicast packets 14824 3255
Broadcast packets 3 0
Multicast packets 6339 3193
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                            16091
Code violations                                0
Filter statistics:
  Input packet count                           9
  Input packet rejects                         9
  Input DA rejects                             9
  Input SA rejects                             0
  Output packet count                          0
  Output packet pad count                     0
  Output packet error count                   0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth          Buffer Priority
Limit
      %          bps          %          usec          low
0 best-effort          95          9500000000          95          0
none
3 network-control          5          500000000          5          0
none
Interface transmit statistics: Disabled
```

Sample Output

**show interfaces
extensive (T4000)**

```
user@host> show interfaces xe-4/0/0 extensive
Physical interface: xe-4/0/0, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 859, Generation: 173
```

Routers with Type 5
FPCs)

```

Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 8 maximum usable queues
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:12:1e:37:53:f8, Hardware address: 00:12:1e:37:53:f8
Last flapped   : 2012-06-06 02:25:56 PDT (10:11:58 ago)
Statistics last cleared: 2012-06-06 12:36:59 PDT (00:00:55 ago)
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, Runt: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort        0                0                0
1 expedited-fo       0                0                0
2 assured-forw       0                0                0
3 network-cont       0                0                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control
Active alarms : None
Active defects : None
PCS statistics      Seconds
Bit errors         0
Errored blocks     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
Packet Forwarding Engine configuration:
Destination slot: 4
CoS information:
Direction : Output
CoS transmit queue          Bandwidth          Buffer Priority
Limit
                                %          bps          %          usec
0 best-effort                95      4750000000    95          0      low
none
3 network-control            5       250000000     5          0      low
none
Interface transmit statistics: Disabled

Logical interface xe-4/0/0.0 (Index 93) (SNMP ifIndex 834) (Generation 158)
Flags: SNMP-Traps 0x4004000 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
Protocol inet, MTU: 1500, Generation: 192, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 34.1.1/24, Local: 34.1.1.2, Broadcast: 34.1.1.255, Generation:
157
Protocol multiservice, MTU: Unlimited, Generation: 193, Route table: 0
Policer: Input: __default_arp_policer__

```

Sample Output

show interfaces
extensive (T4000
Routers with 24-port
10-Gigabit Ethernet

```

user@host> show interfaces xe-3/1/0 extensive
Physical interface: xe-3/1/0, Enabled, Physical link is Up
Interface index: 160, SNMP ifIndex: 1285, Generation: 163
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loopback: None,

```

LAN/WAN PIC on Type 5 FPC)

```

Source filtering: Disabled, Flow control: Enabled
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: 2c:6b:f5:e1:cb:39, Hardware address: 2c:6b:f5:e1:cb:39
Last flapped   : 2012-05-09 07:15:54 UTC (03:39:52 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, Runt: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
  L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0,
  HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0                0                0
  1 ay_q1                    0                0                0
  2 assured-forw             0                0                0
  3 network-cont             0                0                0

Queue number:      Mapped forwarding classes
  0                best-effort
  1                ay_q1
  2                assured-forwarding
  3                network-control
Active alarms   : None
Active defects  : None
PCS statistics          Seconds
  Bit errors              0
  Errored blocks          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
Packet Forwarding Engine configuration:
  Destination slot: 3
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

Preclassifier statistics:
Traffic Class      Received Packets  Transmitted Packets  Dropped Packets

network-control    0                      0                      0
best-effort        0                      0                      0
Interface transmit statistics: Disabled

```

Sample Output

**show interfaces
extensive (Aggregated
Ethernet)**

```

user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 199, SNMP ifIndex: 570, Generation: 202
Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Current address: 2c:6b:f5:d1:0f:c0, Hardware address: 2c:6b:f5:d1:0f:c0
Last flapped   : 2012-06-06 23:33:03 PDT (00:00:58 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          18532          1984 bps
Output bytes :           0           0 bps
Input packets:          158           2 pps
Output packets:           0           0 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :           0
Input packets:           0
Output packets:           0
Dropped traffic statistics due to STP State:
Input bytes :           0
Output bytes :           0
Input packets:           0
Output packets:           0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
Resource errors: 0

```


Output errors:

Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:

0

Ingress queues: 8 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 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: 8 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	57	57	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	63605	63605	0

Queue number:	Mapped forwarding classes
0	best-effort
1	expedited-forwarding
2	assured-forwarding
3	network-control

Logical interface ae0.0 (Index 331) (SNMP ifIndex 583) (Generation 142)

Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2

Statistics	Packets	pps	Bytes	bps	
Bundle:					
Input :	149	2	17416	1984	
Output:	0	0	0	0	
Link:					
ge-3/2/5.0					
Input :	90	1	10100	992	
Output:	0	0	0	0	
ge-3/3/9.0					
Input :	59	1	7316	992	
Output:	0	0	0	0	
LACP info:					
Port	Port	Role	System	System	Port
		priority	identifier	priority	number
key					
ge-3/2/5.0	Actor	100	00:00:00:00:00:01	127	1
1 ge-3/2/5.0	Partner	127	00:24:dc:98:67:c0	127	1 1
ge-3/3/9.0	Actor	100	00:00:00:00:00:01	127	2
1 ge-3/3/9.0	Partner	127	00:24:dc:98:67:c0	127	2 1
LACP Statistics:					
	LACP Rx	LACP Tx	Unknown Rx	Illegal Rx	
ge-3/2/5.0	38	137	0	0	
ge-3/3/9.0	36	139	0	0	
Marker Statistics:					
	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx	
ge-3/2/5.0	0	0	0	0	

```

    ge-3/3/9.0          0          0          0          0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255, Generation:
153
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
  Flags: Is-Primary
  Policar: Input: __default_arp_policer__

```

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 10 on page 212](#). 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 10 on page 212](#).

Table 10: 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:

Framereelay header	: 2 bytes
Control,NLPID	: 2 bytes
Fragmentaion header	: 2 bytes
Inner proto	: 2 bytes
HDLC flag and FCS	: 4 bytes

Fragments 2 ...n :

Framereelay 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
Framereelay header	: 2 bytes
Inner proto	: 2 bytes
HDLC flag and FCS	: 4 bytes

Fragments 2 ...n :

Fragmentaion header	: 2 bytes
Framereelay 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):

=====

Framereelay 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 220](#)
[show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 221](#)

[show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 222](#)
[show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 222](#)
[show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 226](#)
[show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 229](#)
[show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 232](#)
[show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 232](#)
[show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 234](#)
[show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 237](#)
[show interfaces queue \(QFX Series\) on page 247](#)
[show interfaces queue l2-statistics \(lsq interface\) on page 248](#)

Output Fields [Table 11 on page 215](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 11: 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” on page 111 .
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 11: 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. <ul style="list-style-type: none"> • 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 12 on page 218 .
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 211
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by PIC type. For more information, see Table 12 on page 218 . 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 211
Tail-dropped packets	Number of packets dropped because of tail drop.

Table 11: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> (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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type. For more information, see Table 12 on page 218.</p> <ul style="list-style-type: none"> (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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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 12 on page 218](#) shows how the byte counts on the outbound interfaces vary depending on the PIC type. [Table 12 on page 218](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 12: Byte Count by PIC Type

PIC Type	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> 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. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Switches:</p> <ul style="list-style-type: none"> 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. 	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 12: Byte Count by PIC Type (*continued*)

PIC Type	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	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	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>M Series routers:</p> <ul style="list-style-type: none"> 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	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	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	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		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
=====
```

